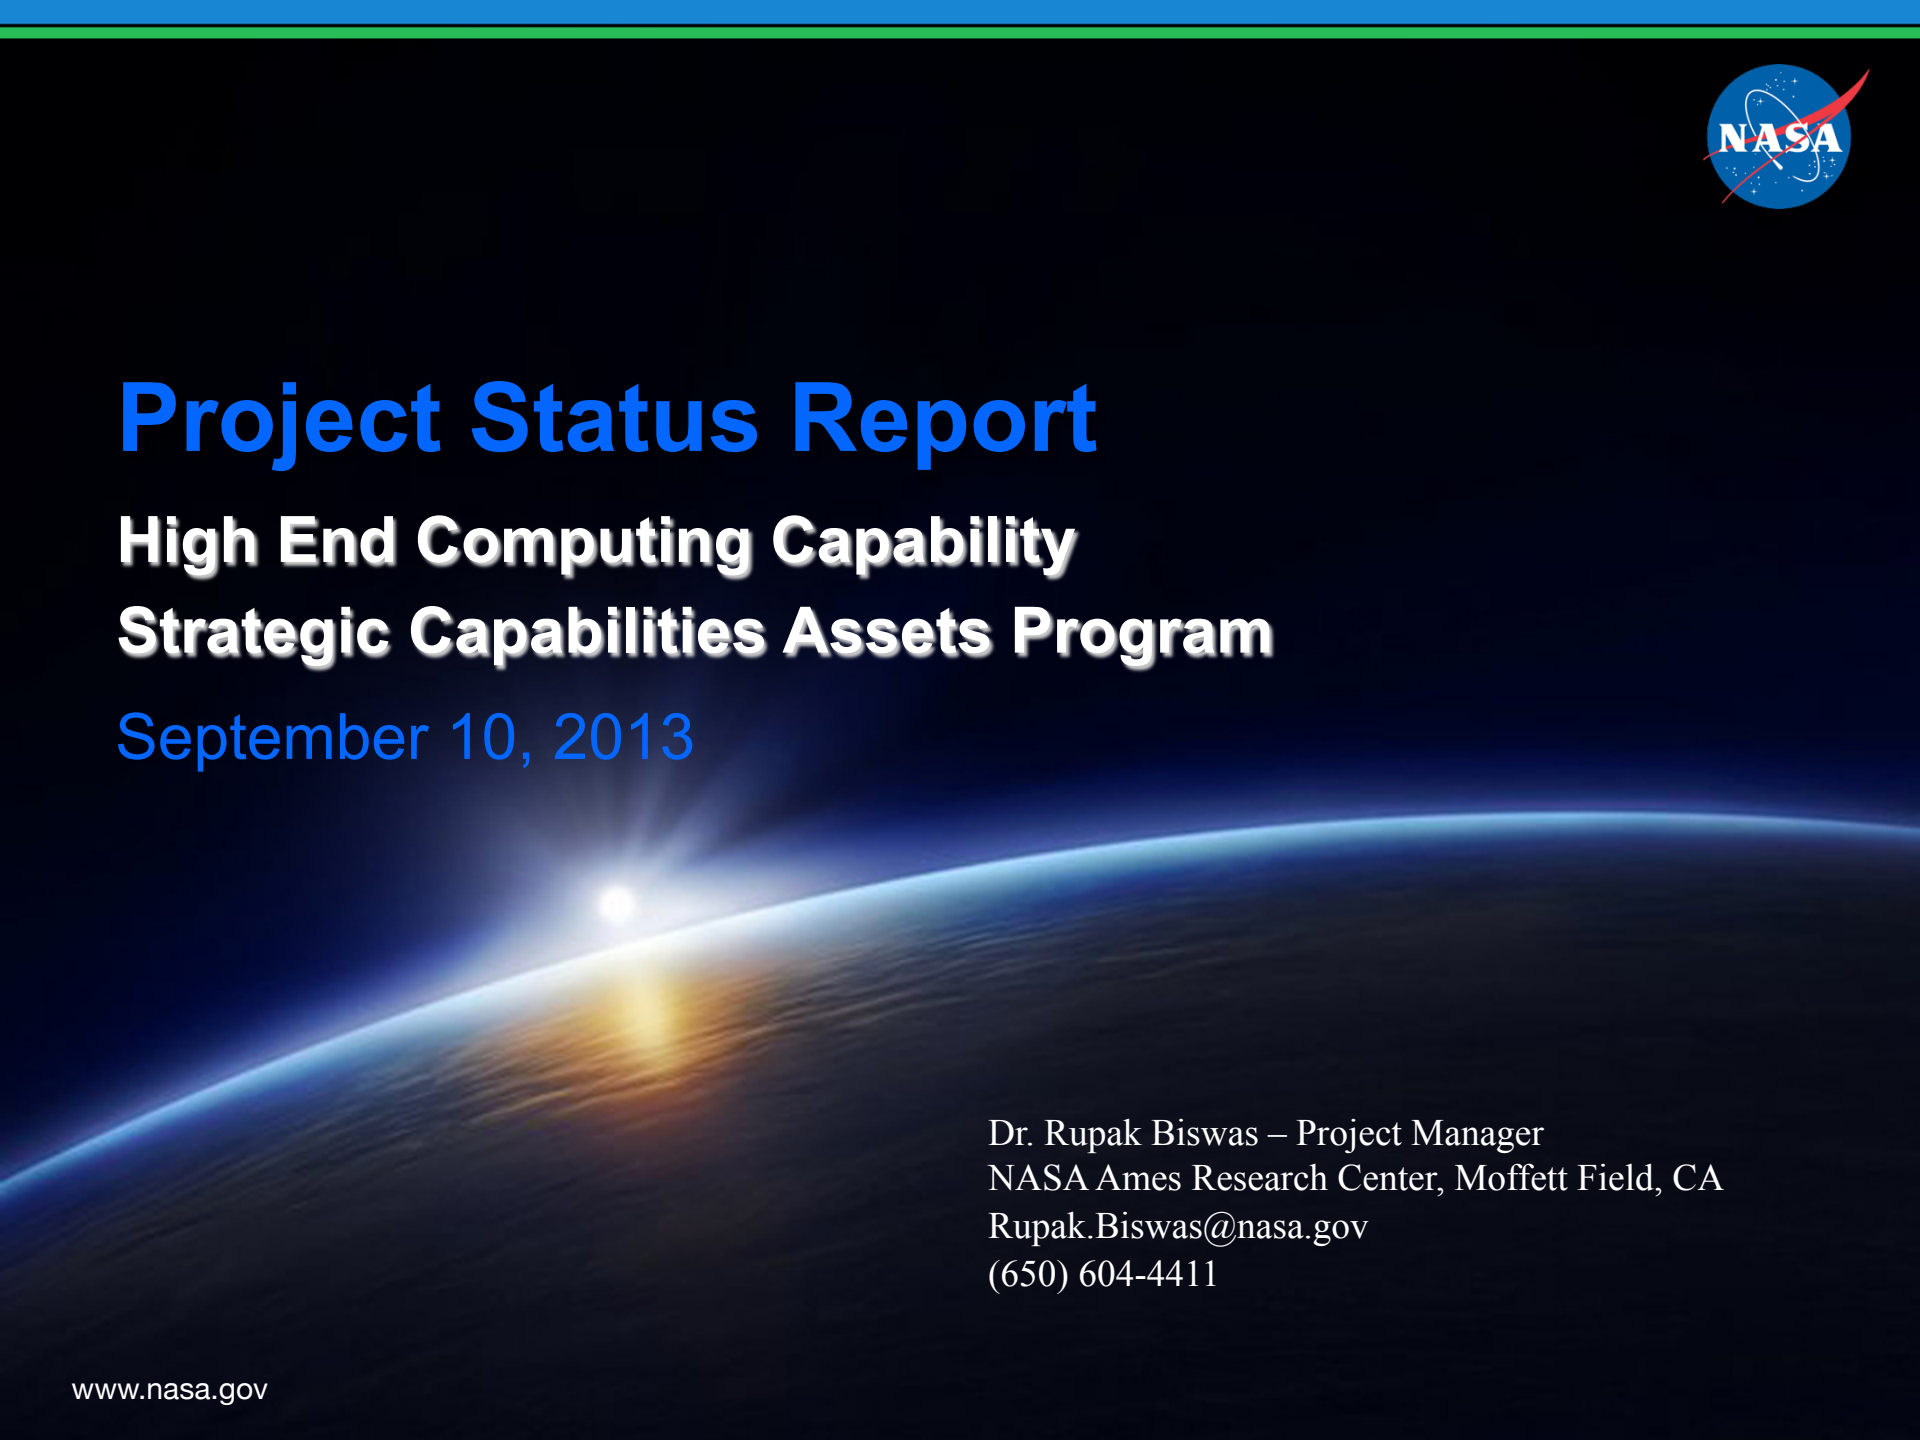

Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

September 10, 2013

Dr. Rupak Biswas – Project Manager
NASA Ames Research Center, Moffett Field, CA
Rupak.Biswas@nasa.gov
(650) 604-4411

Pleiades' Ivy Bridge Upgrade Completed on Schedule

- The HECC Supercomputing Systems team completed the multi-phase Pleiades upgrade on August 24th, with 46 new Ivy Bridge racks replacing 64 Harpertown racks. A subset of the removed Harpertown nodes are being installed on our secondary floor for testing and production. Pleiades' overall peak performance increased nearly 61%—from 1.8 petaflops (PF) to 2.9 PF.
- The new racks provide 1.48 PF of peak performance—3.77 times the peak performance and 4.5 times the SBUs of the Harpertown racks.
- Racks were delivered in three shipments and integrated into Pleiades as soon they were installed by the vendor, to rapidly enable user access.
- The first 16 racks were installed, integrated, tested, and released into production within two weeks of arrival providing more capacity than the removed 64 Harpertown racks. All racks were in production one month after the first shipment arrived.
- The APP and Tech Pubs teams created new and updated existing documentation to assist users with running jobs on the Ivy Bridge nodes.

Mission Impact: Continuous augmentation of the HECC supercomputing environment provides an improved computational capability to support the ever-increasing requirements of NASA's mission directorates.

All 46 of the Ivy Bridge racks were released into production within four weeks after the first of three shipments. Pleiades now has a total of 166 racks containing 11,136 nodes, and the system's overall peak performance has increased to 2.87 petaflops.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Science Corp.

New NFS Filesystem Increases NFS Raw Storage by a Factor of 60

- HECC system engineers released a new nobackup filesystem that employs the Network File System (NFS) protocol.
- The filesystem provides 1 petabyte (PB) of raw storage space; a 60x increase in capacity over the existing NFS nobackup filesystem.
- The new filesystem is configured with a leading-edge feature, known as T10-PI, to provide end-to-end data validation between the NFS server and storage device. This feature provides protection against silent (undetected) data corruption.
- The NFS filesystem is optimized to improve I/O performance for small files, complementing the Lustre nobackup filesystem (see slide 5), which is optimized to improve I/O performance for large files.

Mission Impact: The increased NFS storage capacity enables NASA users with small file I/O access patterns to make more effective use of HECC computing resources.

The new NFS nobackup filesystem has 60 times the storage capacity of the previous NFS filesystem.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Science Corp.

HECC Doubles Lustre Storage Capacity for Compute Systems

- The HECC Supercomputing Systems team deployed two new Lustre filesystems for the HECC computing environment.
- The Lustre filesystems provide a total of 8.5 petabytes (PB) of usable space (11 PB raw) that will double the storage capacity for HECC users in all Mission Directorates.
- The filesystems are configured with a leading-edge feature, known as T10-PI, to provide end-to-end data validation. This feature provides protection against silent (undetected) data corruption.
- The filesystems have been optimized to improve random I/O performance, compared to the existing filesystem, and will result in better interactive responsiveness for HECC users.
- Users will be migrated to the filesystem over the next several months to minimize the impact on their workflow.

Mission Impact: The significant increase in storage capacity and functionality will enable NASA researchers to more fully utilize HECC computing resources and run more data-intensive applications.

HECC's Lustre storage will be split between two systems, one with 1.8 petabytes (PB) capacity and the other with 6.7 PB capacity.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Science Corp.

Archive Data Successfully Recovered from Corrupted Tape Media

- HECC storage experts successfully recovered 15 terabytes (TB) of archived user data from corrupted LTO-5 tape media, after an event that resulted in 16 tapes having an “End-of-Data Marker” placed at the beginning of each tape.
- Working with SGI and Spectra Logic (the tape library manufacturer), the HECC team determined the issue was caused by an incorrect operating system (OS) behavior. The vendor has corrected this flaw.
- HECC’s data retention policy requires writing two copies of the data: one to the primary NAS facility, and one to the secondary facility located about a kilometer away. Writing two copies of data enables HECC to recover from most data loss events.
- Due to this policy, only 5 tapes (totaling 142 gigabytes of data) were unrecoverable using standard processes available to HECC.
- Spectra Logic staff worked with their tape drive supplier (IBM) to recover data on the remaining 5 tapes. This saved NASA about $50K in fees that would have been charged by a commercial data recovery service.

Mission Impact: HECC’s storage expertise, collaborative relationship with vendors, and the policy of retaining two copies of data, enables the project to minimize the risk of data loss for NASA users, and saves significant cost to the agency.

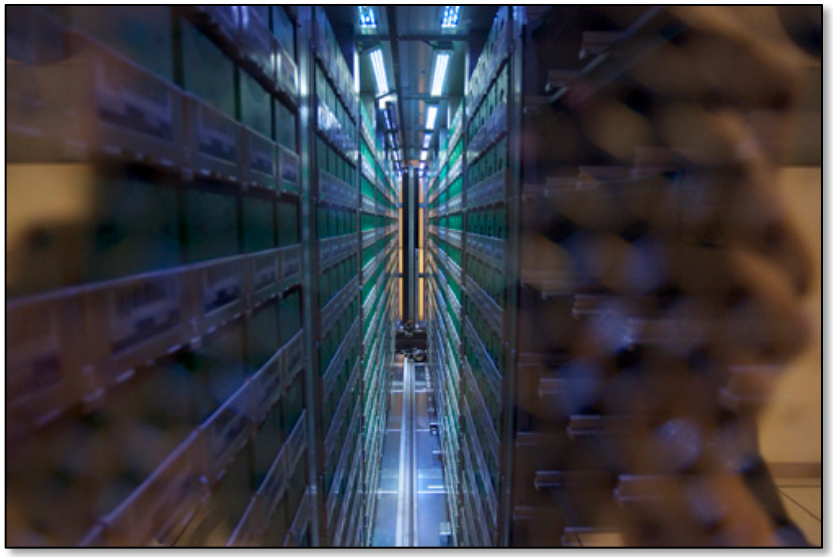

Each LTO-5 tape in HECC’s tape library can hold 1.5 terabytes of uncompressed data. Collectively, these tapes would stretch more than 30,273 miles—long enough to completely wrap around Earth at the equator.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Science Corp.

Record-High Usage on Pleiades Provides Maximum Resources for NASA Missions

- August showed record-high usage of the Pleiades supercomputer, with 8.27 million Standard Billing Units (SBUs) used by NASA's science and engineering organizations—exceeding the previous record of 6.867 million SBUs by over 20%.
- This increase was enabled by the addition of 3,312 Ivy Bridge nodes, and through efficient operations that delivered better than 96% availability and over 83% system utilization (target utilization is 75%).
- The Kepler Science Operations Center Pipeline used almost 900,000 SBUs, accounting for 11% of total usage.
- The next nine projects (4 ARMD, 2 HEOMD, 1 NESC, 2 SMD) each used more than 170,000 SBUs and accounted for another 32% of the usage.
- Computing resources continue to expand to keep pace with the increasing needs of users from all Mission Directorates.

Mission Impact: Increasing Pleiades' system capacity provides mission directorates with more resources for the accomplishment of their goals and objectives.

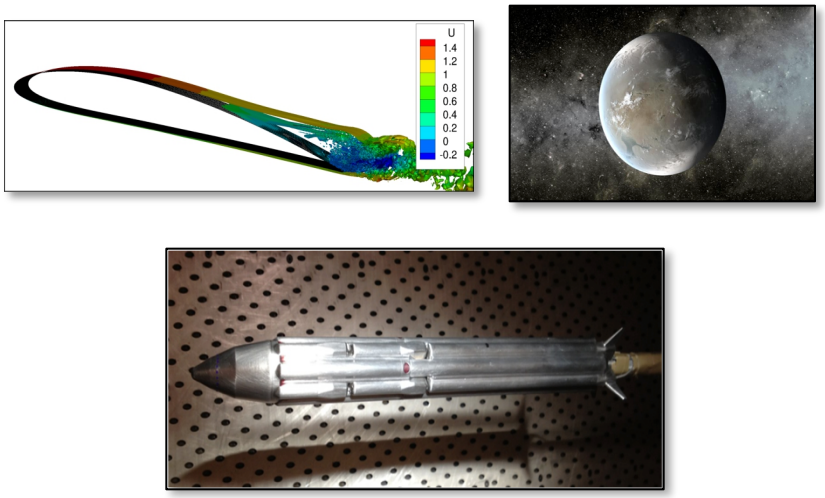

Images representing projects that used the most time in their respective Mission Directorates. Clockwise from top left: Instantaneous vortical structures on the NACA 4412 airfoil, ARMD, K. Shariff; artist's concept of Kepler-62f, a super-Earth-size planet in the habitable zone of a star smaller and cooler than the sun, SMD, T. Klaus; Image of the SWORDS launch vehicle model during installation for a wind tunnel test, HEOMD, S. Krist.

POC: Blaise Hartman, blaise.hartman@nasa.gov, (650) 604-2539, NASA Advanced Supercomputing Division, Computer Sciences Corp.

HECC Team Enhances Network Flow Analysis Capability with Visualization Tool

- HECC network engineers deployed a new server to collect both metadata and statistical information on users' data transfer flow characteristics, with Argus software.
- Data collected allows engineers to visualize historical trends and easily identify when some activity or event is impacting network performance for a particular system or site.
- Critical details about flow characteristics such as packet loss, retransmit counts, and TCP window size in each direction show whether a performance issue is consistent or intermittent, asymmetric or bi-directional, and whether it impacts an entire site or just a single system.
- This new capability enables visualization of a large number of data points, with the ability to zoom into specific problem areas and pinpoint exact timeframes of problem events.

Mission Impact: Tools developed to aid in quick and proactive identification of network problems reduce troubleshooting time and enable faster problem resolution for HECC users.

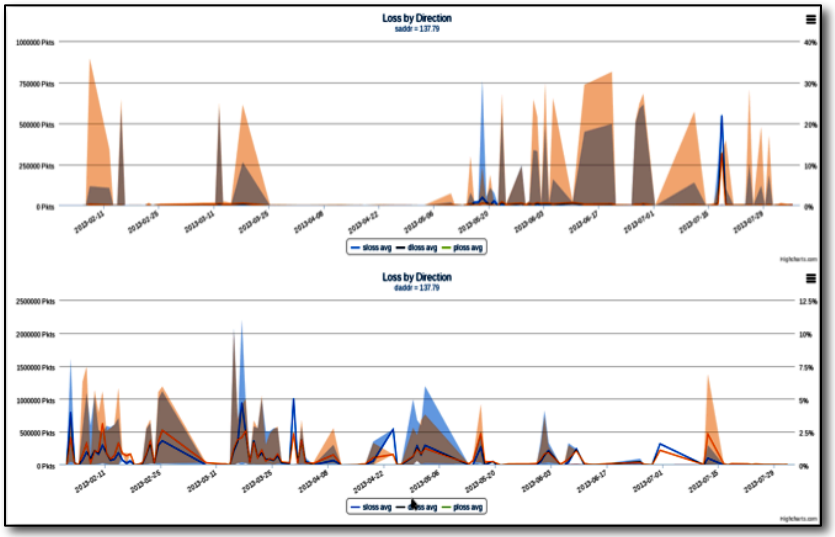

This graph shows the maximum, minimum, and average number of packet drops per day for both source and destination traffic from a Jet Propulsion Laboratory network known to have issues.

POC: Nichole Boscia, nichole.boscia@nasa.gov, (650) 604-0891, NASA Advanced Supercomputing Division, Computer Sciences Corp.

Network Team Evaluates 40-Gigabit Ethernet and RDMA Technology

- Network engineers used Mellanox Network Interface Cards (NICs) and a demo version of the Cisco Nexus 40-gigabit (40G) switch to determine the potential for 40G Ethernet technology to increase bandwidth to HECC servers and front-end systems.
- The team ran single-stream throughput tests between two end hosts using both traditional Ethernet and Remote Direct Memory Access (RDMA) over Ethernet (RoCE) technologies.
- The team optimized two test systems and 40G NICs to identify configurations that gave the best sustained rates for each technology.
  - RoCE single-stream TCP session achieved a sustained rate of 39 gigabits per second (Gbps).
  - Traditional Ethernet single-stream TCP session achieved a sustained rate of 21 Gbps.
- The team also produced a whitepaper detailing the setup, tuning, and test results obtained during this evaluation.

Mission Impact: The next-generation Ethernet will enable a 4x improvement in data throughput over existing infrastructure, and can leverage direct-memory capabilities to converge networking technologies in data centers like the HECC secure enclave. Early testing on future environments identifies the optimum applications to exploit the technology.

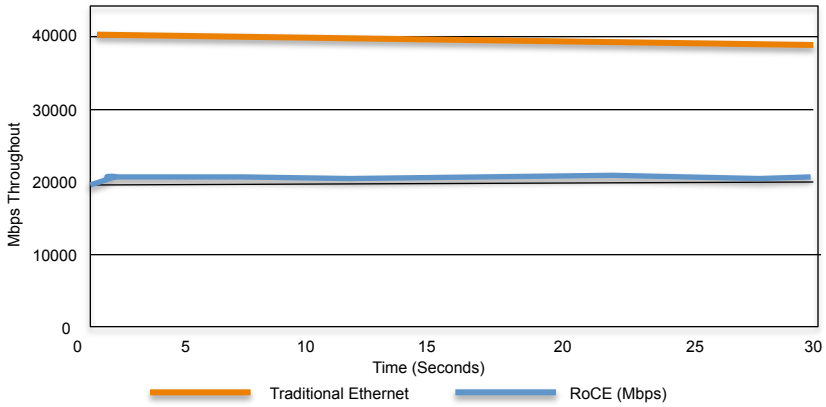

This graph shows a performance comparison between 40G traditional Ethernet and 40G using Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) technologies, using single-stream TCP testing. The RoCE performance was almost twice as fast compared to traditional Ethernet

POCs: Nichole Boscia, nichole.boscia@nasa.gov, (650) 604-0891, Harjot Sidhu, harjot.s.sidhu@nasa.gov, (650) 604-4935, NASA Advanced Supercomputing Division, Computer Sciences Corp.

HECC Resources Enable Extensive CFD Analyses for SLS SRB Separation Database

- Modeling and simulation experts in the NASA Advanced Supercomputing (NAS) Division completed extensive computational fluid dynamics (CFD) simulations of solid rocket booster (SRB) separation for the Space Launch System (SLS).
  - Performed over 2,000 SRB separation simulations using the Cart3D, OVERFLOW, and LAVA CFD codes.
  - Modeled cases covering a wide range of separation distances and angles, with and without plumes from the rocket engines and separation motors.
  - Computed aerodynamic forces, moments, surface pressures, and line loads on the SRBs and core stage.
- These simulations produced valuable aerodynamic data for the SLS SRB Separation Database.
  - Generated the full set of 777 database cases using the efficient, inviscid Cart3D code.
  - Performed viscous OVERFLOW cases to verify the inviscid results, establish uncertainty levels, and identify conditions requiring higher-fidelity analysis.
  - Performed additional analyses to establish simulation best practices, assess key sensitivities, etc.
- These extensive analyses required a total of 1.2 million processor-hours on the Pleiades supercomputer since January 2013.

** HECC provided supercomputing resources and services in support of this work*

Mission Impact: Enabled by HECC resources, CFD analyses of SRB separation provide critical aerodynamic data needed to design safe, effective separation systems and assess potential risks of recontact with the core stage.

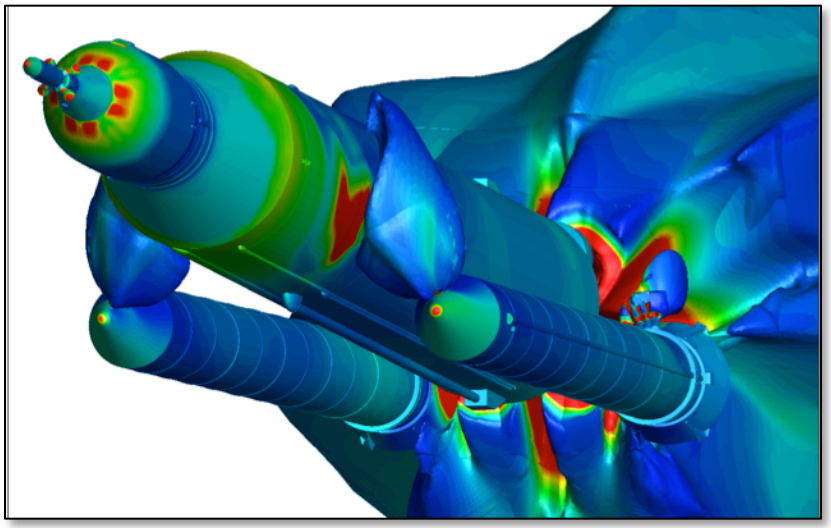

Visualization of surface pressures and jet plume iso-contours from an OVERFLOW stage separation simulation. Jeff Onufer, NASA/Ames

POC: Cetin Kiris, cetin.c.kiris@nasa.gov, (650) 604-4485,
NASA Advanced Supercomputing Division

HECC Resources Enable Quick Turnaround on SLS Valve Design Cycles

- Researchers at NASA Marshall Space Flight Center (MSFC) completed 3D, time-accurate computational fluid dynamics (CFD) simulations on the Pleiades supercomputer, providing valuable input into the Space Launch System valve design.
  - Used the Loci/CHEM CFD program to model the opening and closing cycle of a candidate SLS Flow Control Valve (FCV) to evaluate potentially damaging unsteady fluid environments in the valve poppet, and other critical performance metrics.
  - Through CFD analyses, identified fluid mode shapes and amplitudes that could have caused the FCV to fail, and showed that conditions leading to the failure occurred at some point during valve actuation.
  - Analyses resulted in design modifications made to avoid creating potentially unsteady fluid modes in the system during valve operation.
- These large scale simulations—completed in just a few days on Pleiades—provided essential information about all phases of valve operation prior to manufacturing or testing, and influenced the final design of the valve used in the SLS liquid engine.

Mission Impact: HECC's expanded supercomputing resources enabled Space Launch System design support teams to complete simulations quickly enough to support iterative valve design cycles.

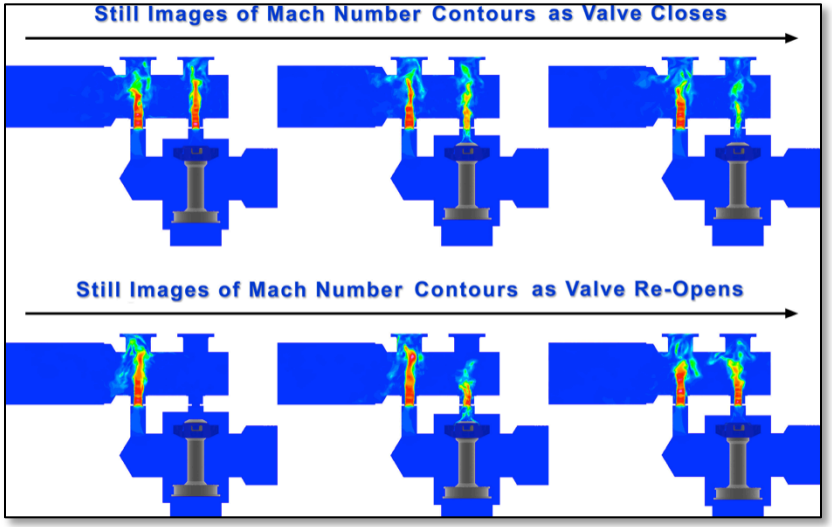

Velocity magnitude snapshots from a moving valve simulation. The top images show valve closure from left to right. The bottom images show the valve reopening from left to right.

** HECC provided supercomputing resources and services in support of this work*

POC: Brian R. Richardson, brian.r.richardson@nasa.gov, (256) 544-9396, NASA Marshall Space Flight Center

Pleiades Plays Critical Role in NASA Earth Exchange (NEX) Collaborations

- Enabled by the Pleiades supercomputer, members of the NASA Earth Exchange (NEX) have delivered several important studies over the last year, highlighting the significance of long-term observations for improving our understanding of the Earth system.
  - Researchers derived a new global vegetation map from over 80,000 NASA Landsat scenes, at 30-meter resolution, to provide a dynamic view of changing vegetation over the course of the year, delivering an unprecedented monitoring capability at a global scale.
  - Using 9 ecosystem models, an international research team found that tropical ecosystems can generate significant carbon dioxide when temperatures rise, unlike ecosystems in other parts of the world.
  - Another NEX team used sophisticated machine learning algorithms to process one-meter, multispectral data covering California to identify and count the number of trees in the state.
- Along with Pleiades, HECC's data storage and high-speed networks enable NEX to provide the capabilities to execute modeling and data analysis on a scale not previously achievable by most scientists.

** HECC provided supercomputing resources and services in support of this work*

Mission Impact: Enabled by HECC's Pleiades supercomputer and other HECC resources, the NASA Earth Exchange (NEX) team is working to lower the barrier of entry to data-and compute-intensive research projects, and to engage global science communities to address grand challenge problems in Earth sciences.

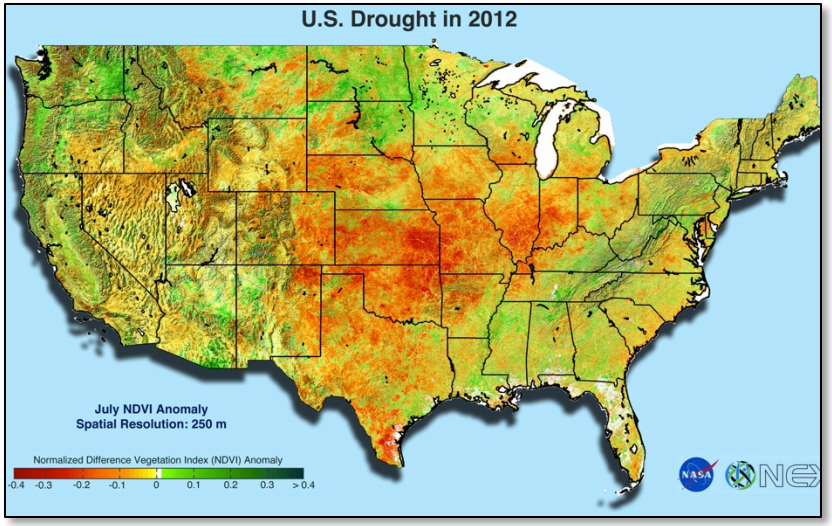

Accelerating analysis with NEX: Rapid assessment of the impacts of 2012 summer drought in the United States on regional ecosystems.

POCs: Petr Votava, petr.votava@nasa.gov, (650) 604-4675,
Ramakrishna Nemani, rama.nemani@nasa.gov. (650) 604-3625,
NASA Ames Research Center

HECC Facility Hosts Several Visitors and Tours in August 2013

- HECC hosted 8 tour groups in August; guests learned about the agency-wide missions being supported by Pleiades, and viewed scientific results on the hyperwall system. Visitors this month included:
  - Mary Lou Jepsen and Carlin Vieri from the Google X Project office met with HECC visualization experts.
  - A group from the Korean Aerospace Research Institute (KARI) received an informational overview and tour of the NAS facility during their visit to Ames.
  - Senior staff from D-Wave Systems and In-Q-Tel met with Ames senior staff and received a tour of the facility, including the quantum computer room.
  - Winners from the Littelfuse Speed2Design contest, hosted by NASA Tech Briefs staff, visited Ames for informal overviews of the center, including a visit to the facility.
  - Staff from Rogue Wave, creator of TotalView software, presented a technical seminar at Ames and received an informal tour of the facility.
  - HECC/NAS staff who received NASA Honor Awards invited family and friends to learn more about their work; guests received informal hyperwall-2 and computer demos after the ceremony.

Winners from the Littelfuse Speed2Design contest viewed scientific visualizations on the hyperwall-2 and received an overview of the HECC Project and the NAS Division.

POC: Gina Morello, gina.f.morello@nasa.gov, (650) 604-4462, NASA Advanced Supercomputer Division

Papers and Presentations

- "**Numerical Simulations of Multiple Scattering of the F-Mode by Flux Tubes**," T. Felipe, A. D. Crouch, A. C. Birch, arXiv:1308.1139 [astro-ph.SR], August 5, 2013. * *http://arxiv.org/abs/1308.1139*
- "**The Acceleration of Thermal Protons at Parallel Collisionless Shocks: Three-Dimensional Hybrid Simulations**," F. Guo, J. Giacalone, The Astrophysical Journal, vol. 773, no. 2, August 5, 2013. * *http://iopscience.iop.org/0004-637X/773/2/158/article*
- "**The Distribution of Dark Matter in the Milky Way's Disk**," M. Kuhlen, A. Pillepich, J. Guedes, P. Madau, arXiv:1308.1703 [astro-ph.GA], August 7, 2013. * *http://arxiv.org/abs/1308.1703*
- "**Confronting Simulations of Optically Thick Gas in Massive Halos with Observations at z=2-3**," M. Fumagalli, J. F. Hennawi, X. Prochaska, D. Kasen, A. Dekel, D. Ceverino, J. Primack, arXiv:1308.1669 [astro-ph.CO], August 7, 2013. * *http://arxiv.org/abs/1308.1669*
- "**Tentative Identification of the Source of Four UHECRs and Implications Thereof**," G. R. Farrar, arXiv:1308.3828 [astro-ph.HE], August 18, 2013. * *http://arxiv.org/abs/1308.3828*

** HECC provided supercomputing resources and services in support of this work*

News and Events

- **The Top 10 Reasons for the US to Return to Space Exploration**, *EDN Network,* August 9, 2013—EDN's Steve Taranovich lists the top ten reasons NASA should continue to be supported, and lists supercomputing—specifically the Pleiades supercomputer—at number five on the list.
  http://www.edn.com/design/analog/4419551/1/The-top-10-reasons-for-the-US-to-return-to-space-exploration
- **How NASA Has Furthered Innovation**, *SFGate*, August 18, 2013—SFGate's James Temple looks at NASA's current technological and scientific innovations, including the 128-screen hyperwall-2 at the NASA Advanced Supercomputing facility, during a press tour of NASA's Ames Research Center.
  http://www.sfgate.com/technology/dotcommentary/article/How-NASA-has-furthered-innovation-4740787.php
- **NASA Visualizations Hit the Hyperwall**, *HPCwire*, August 26, 2013—During a press tour of the NASA Advanced Supercomputing (NAS) facility, Bryan Biegel, NAS deputy division chief, showcases visualizations in aeronautics, space operations, and Earth and space science on the hyperwall-2.
  http://www.hpcwire.com/hpcwire/2013-08-26/nasa_visualizations_hit_the_hyperwall.html
  - **Video: Simulating the Universe with NASA's Supercomputers**, *Design News,* August 26, 2013
    http://www.designnews.com/author.asp?section_id=1395&doc_id=267002

HECC Utilization

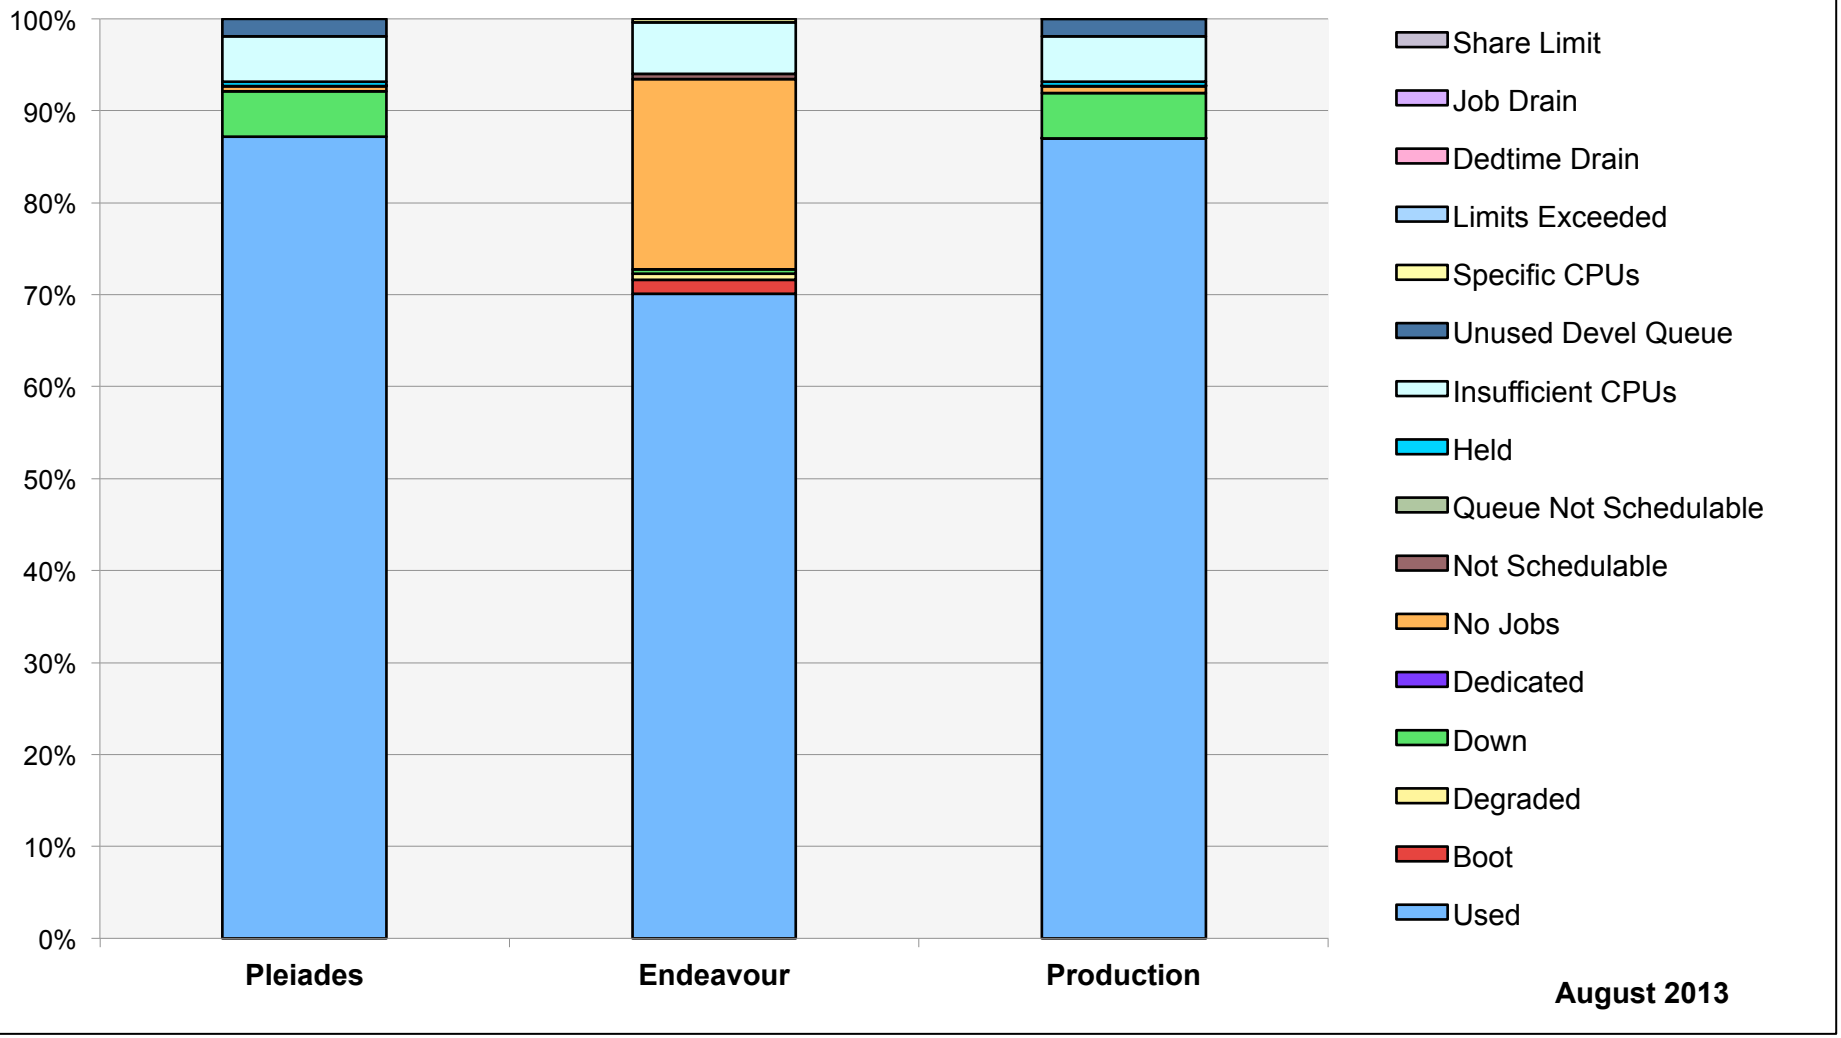

HECC Utilization Normalized to 30-Day Month

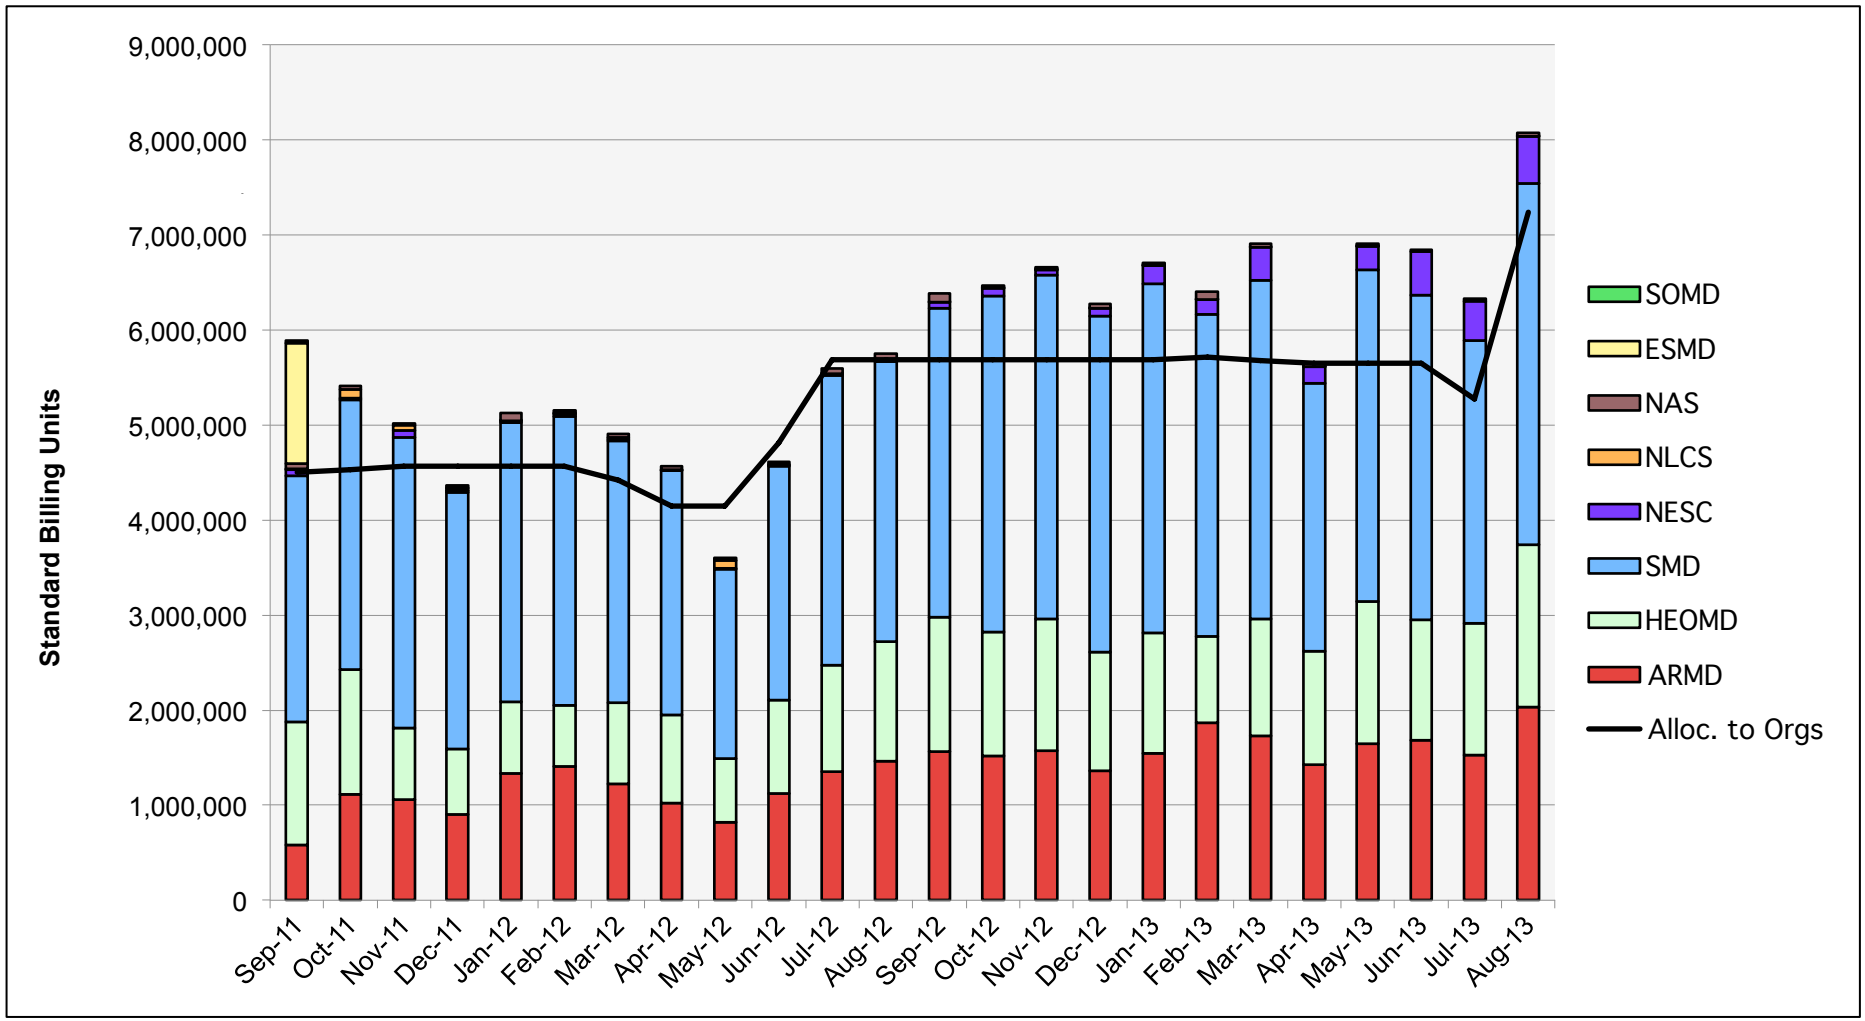

HECC Utilization Normalized to 30-Day Month

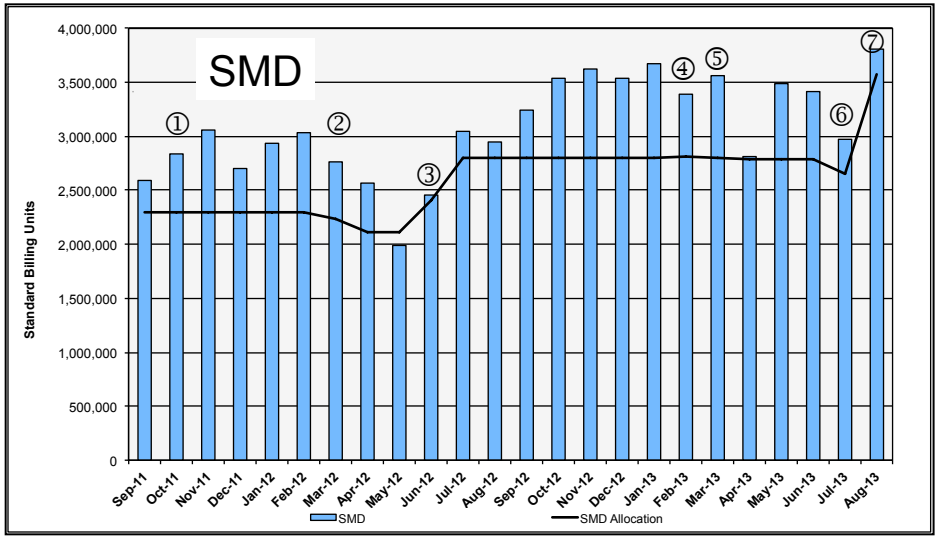

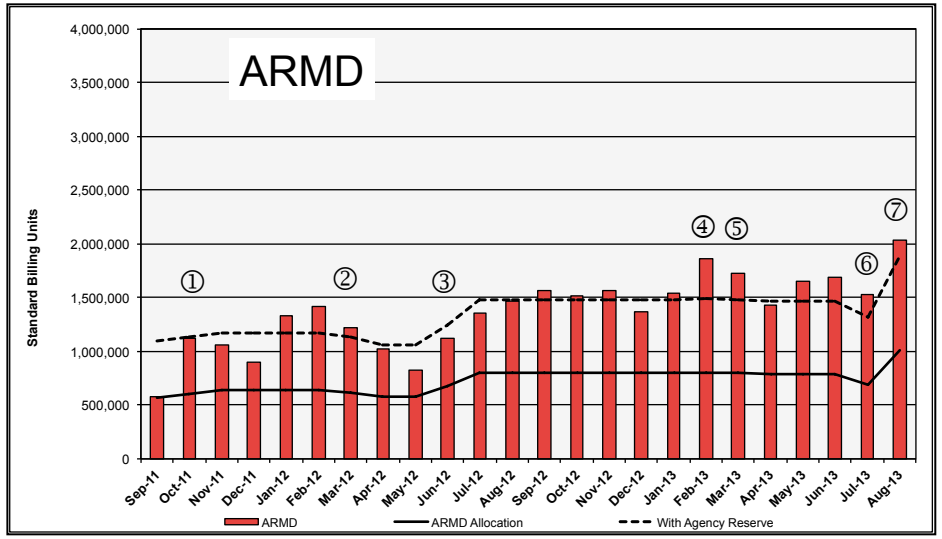

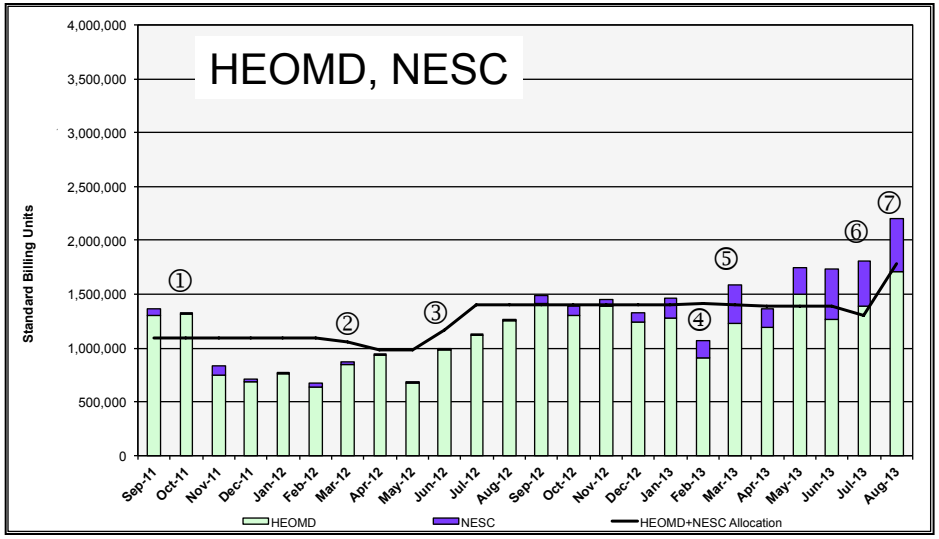

① ARMD Westmere Racks added
② 28 Harpertown Racks removed
③ 24 Sandy Bridge Racks added
④ Columbia 21, 23, and 24 retired; Endeavour 2 added
⑤ Columbia 22 retired; Endeavour 1 added
⑥ 32 Harpertown Racks retired
⑦ 32 Harpertown Racks retired; 46 Ivy Bridge Racks added

Tape Archive Status

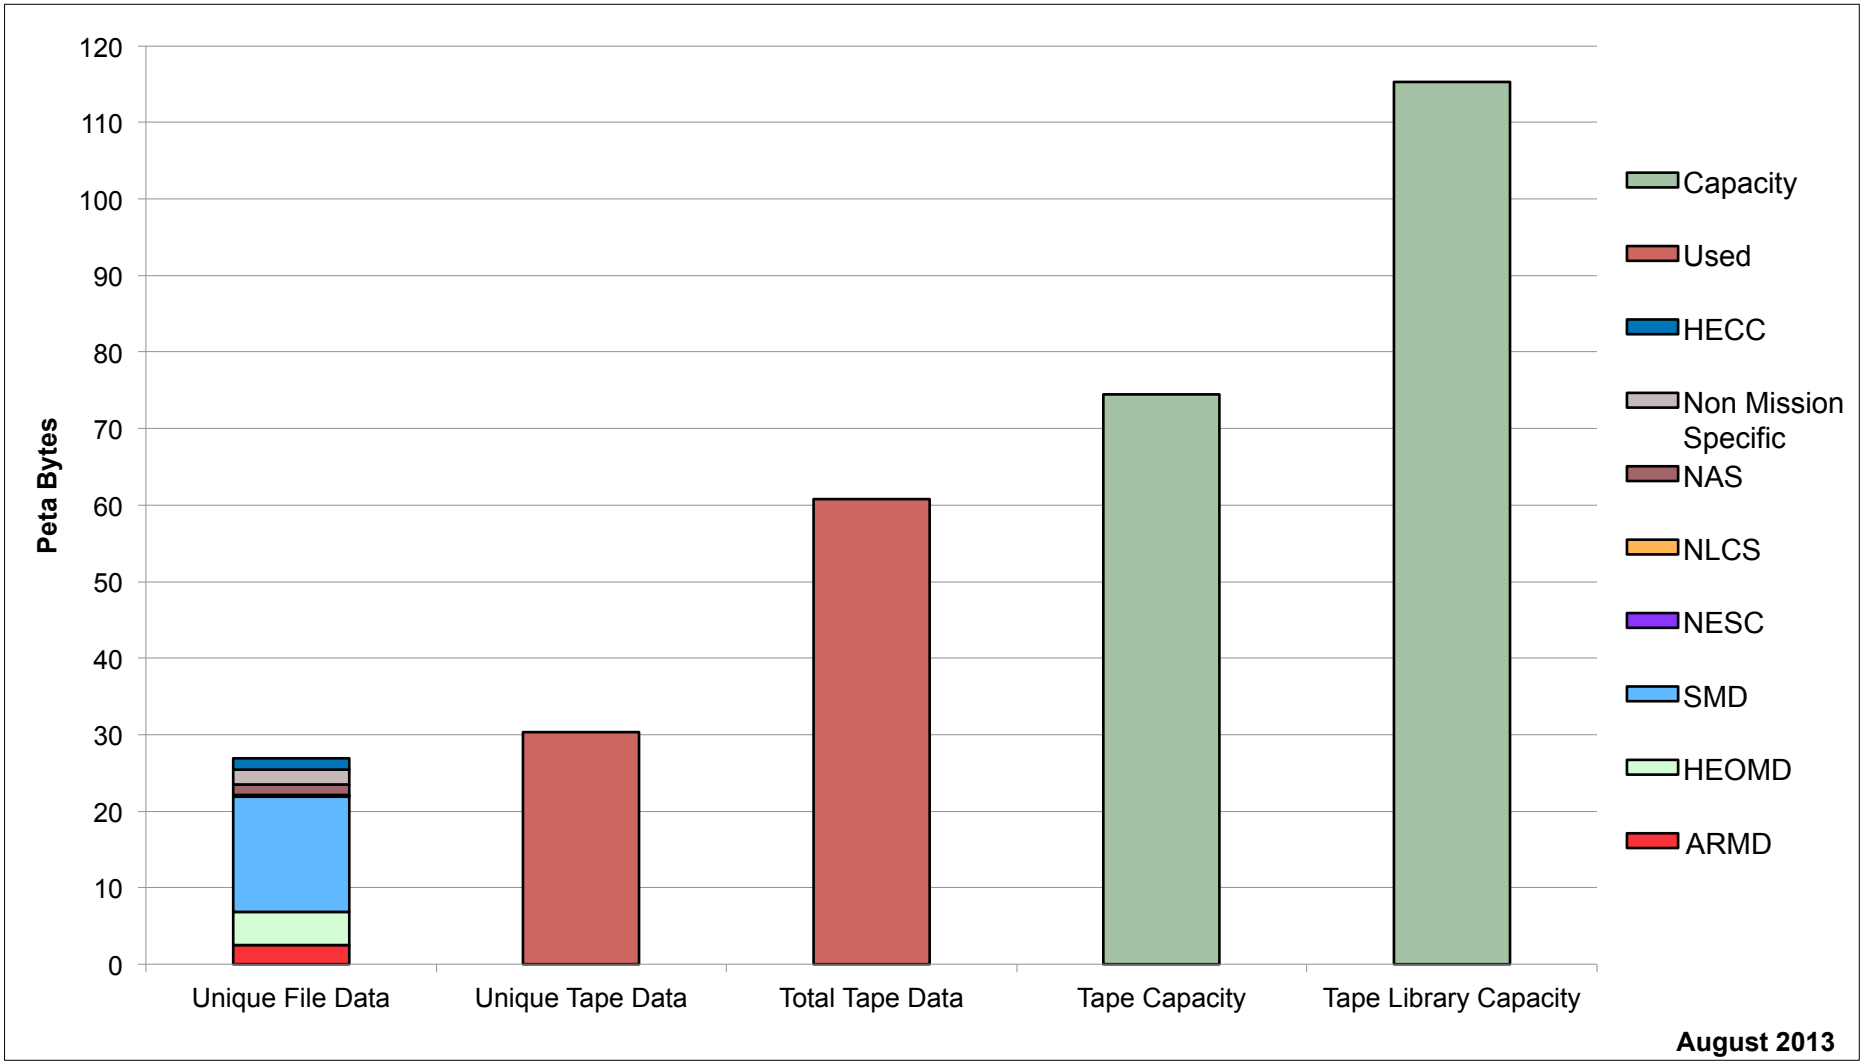

Tape Archive Status

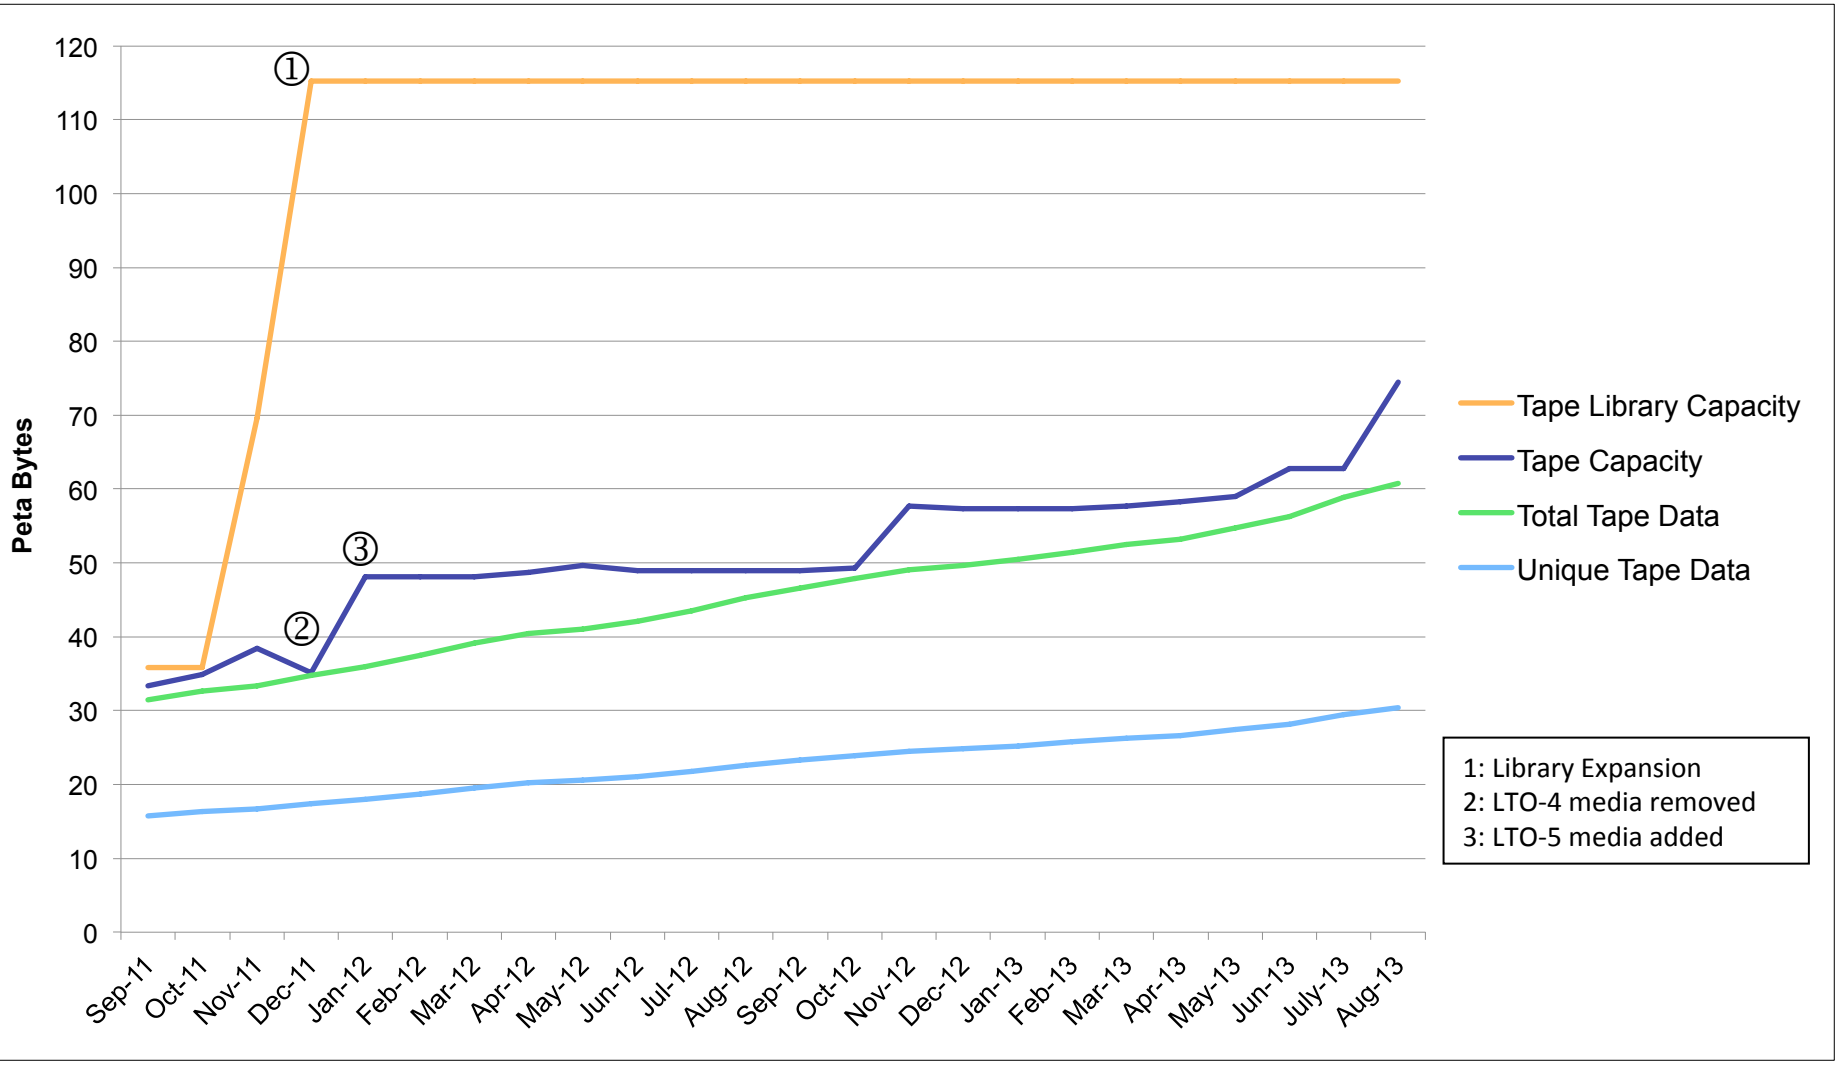

Pleiades:
SBUs Reported, Normalized to 30-Day Month

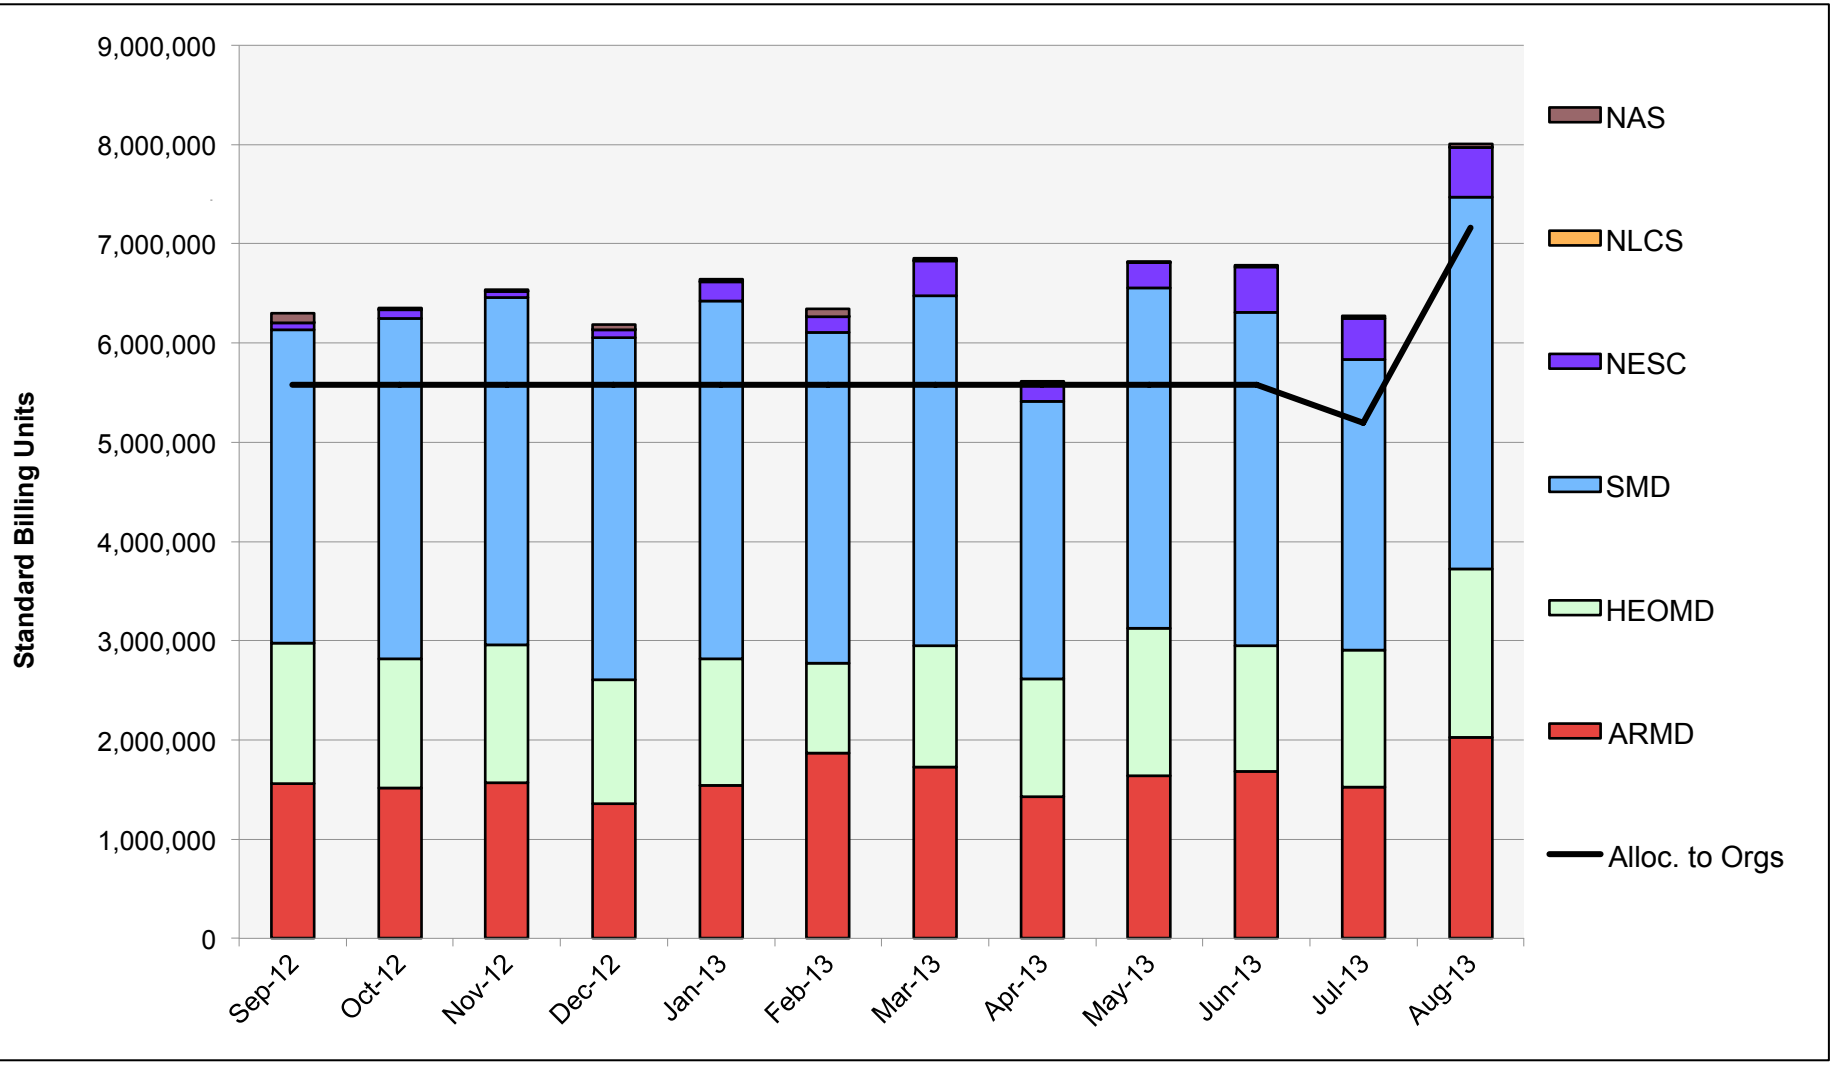

Pleiades:
Devel Queue Utilization

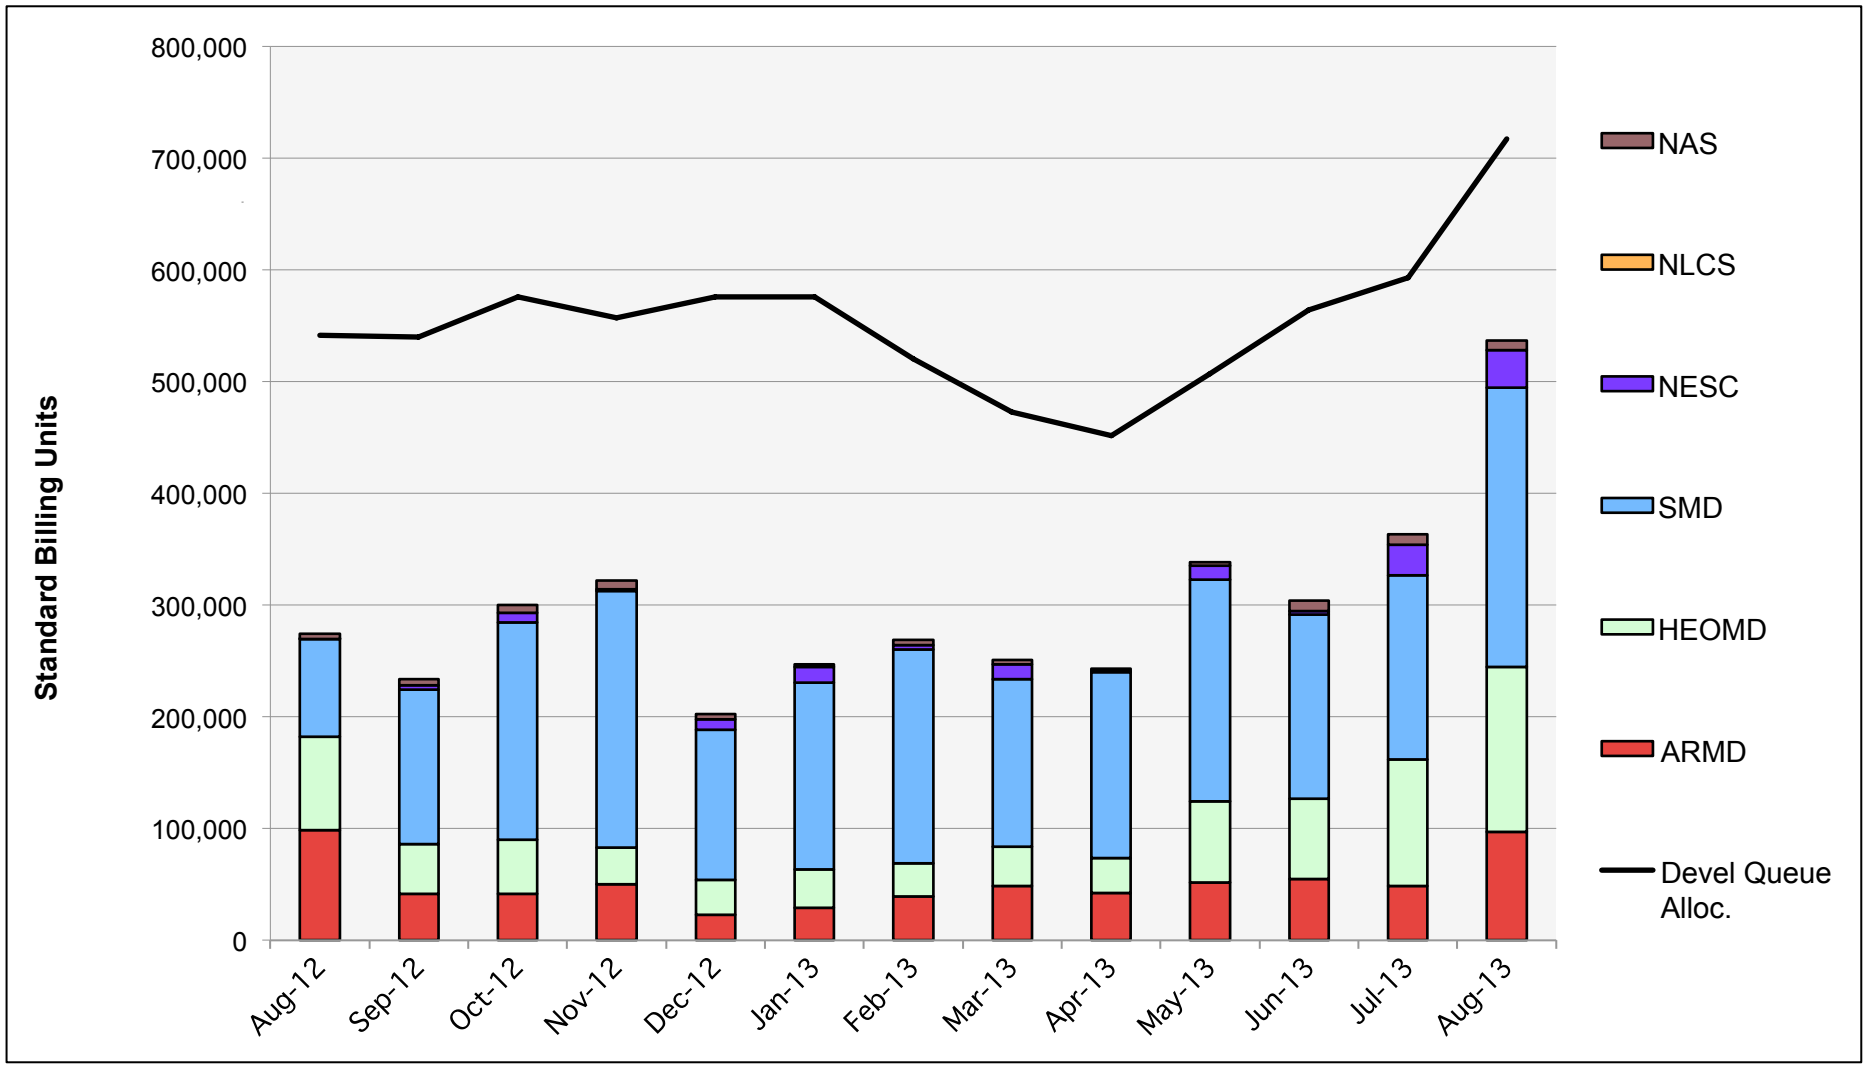

Pleiades:
Monthly Utilization by Job Length

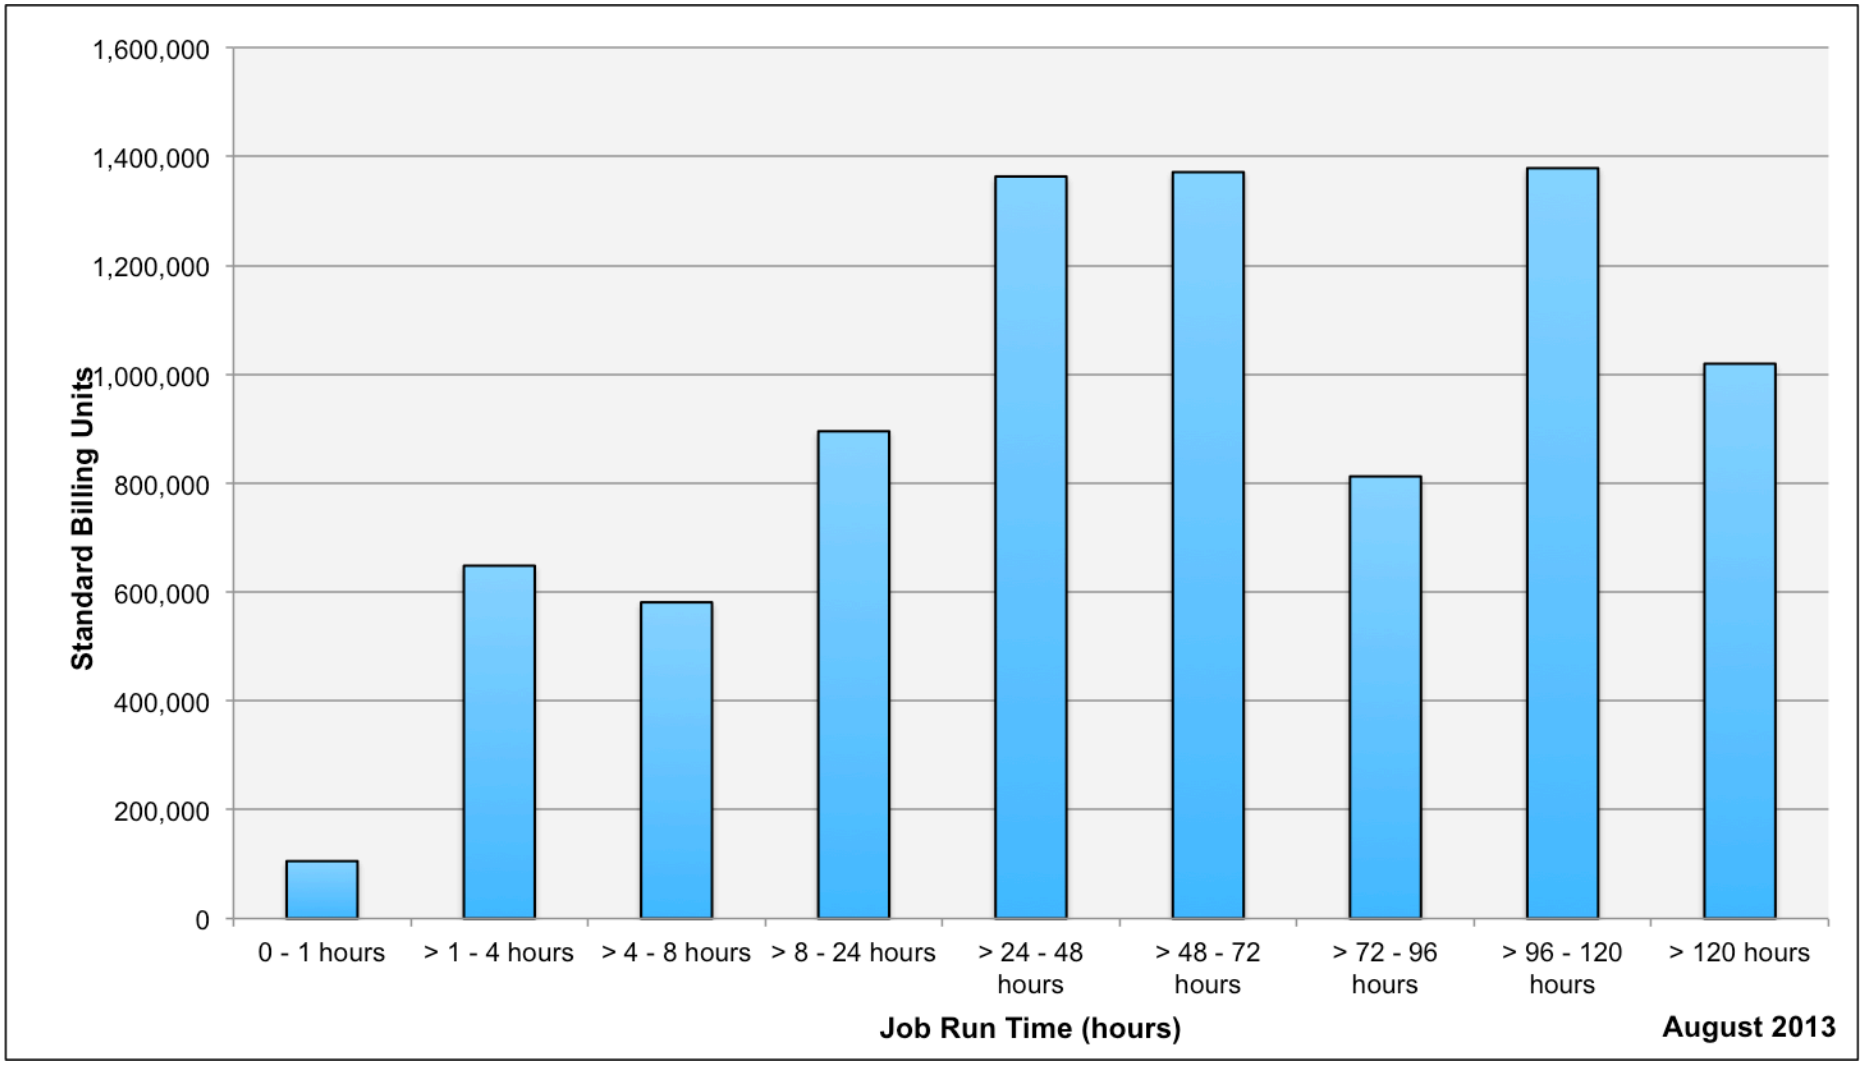

Pleiades:
Monthly Utilization by Size and Mission

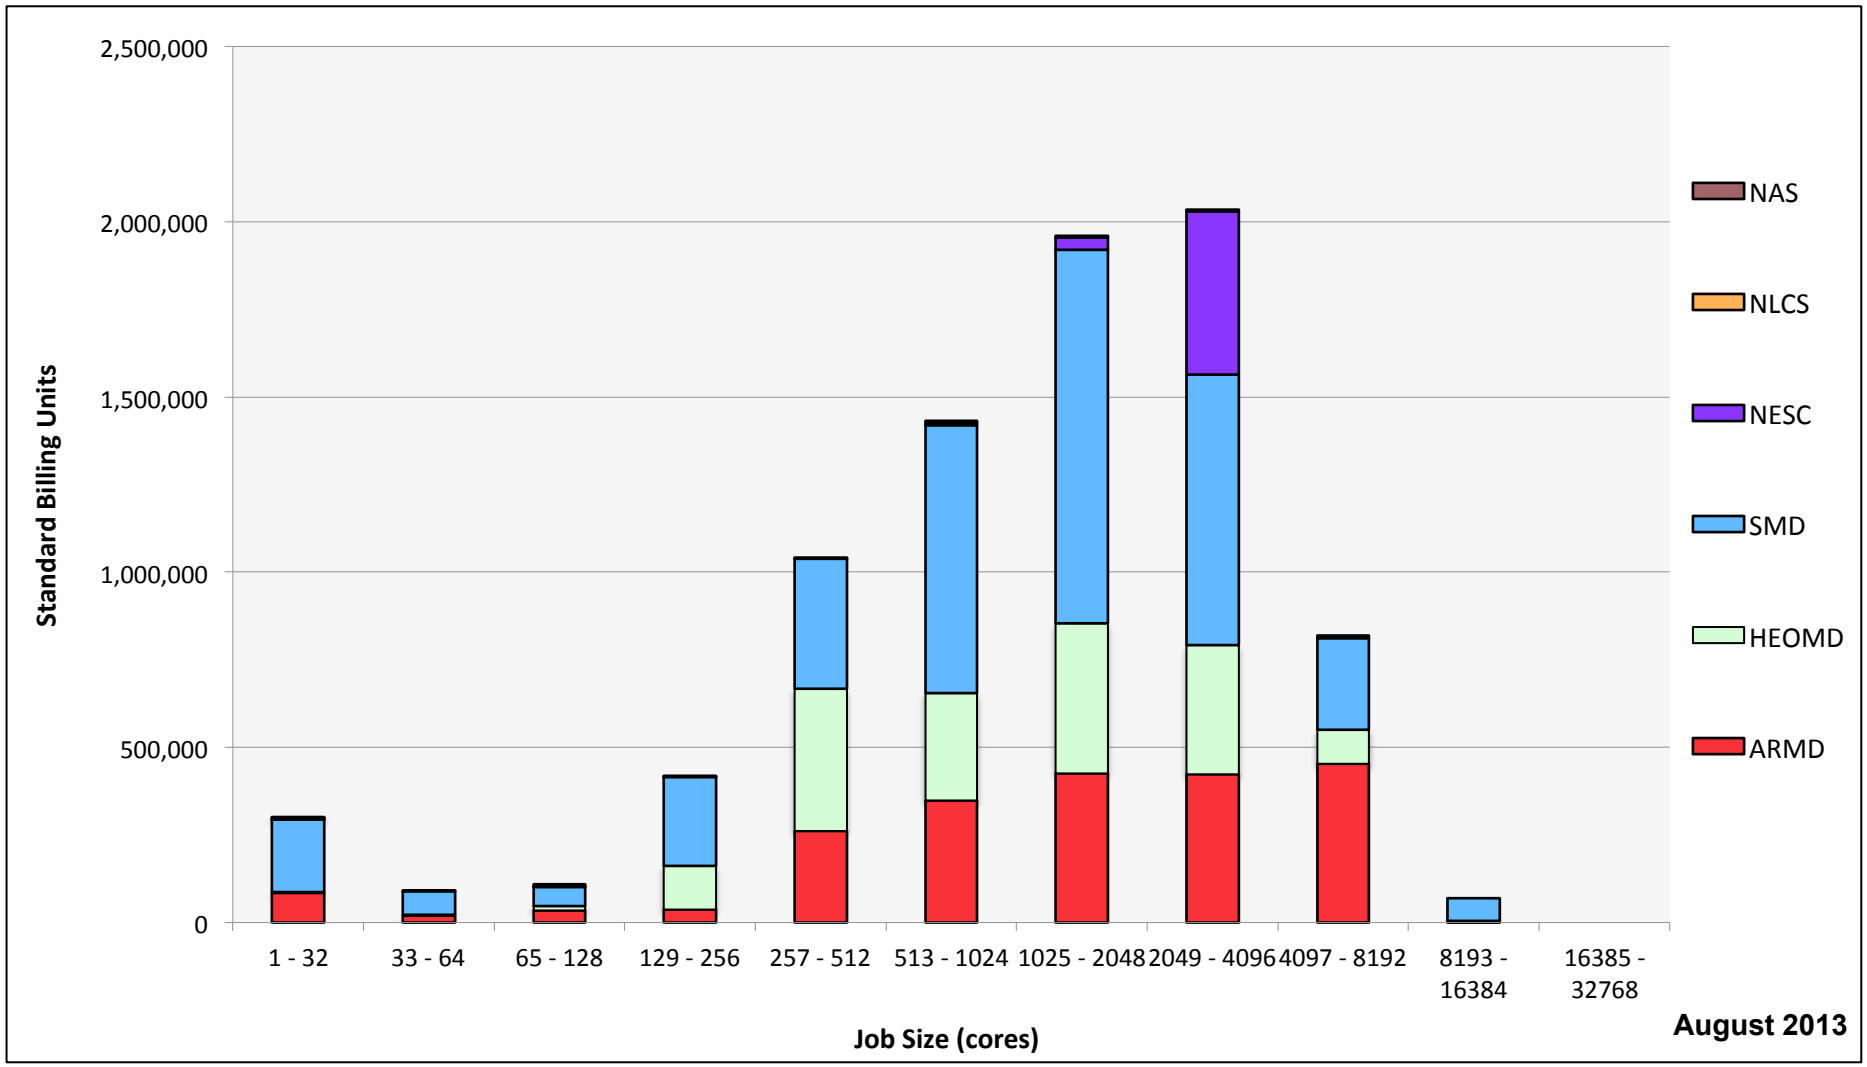

Pleiades:
Monthly Utilization by Size and Length

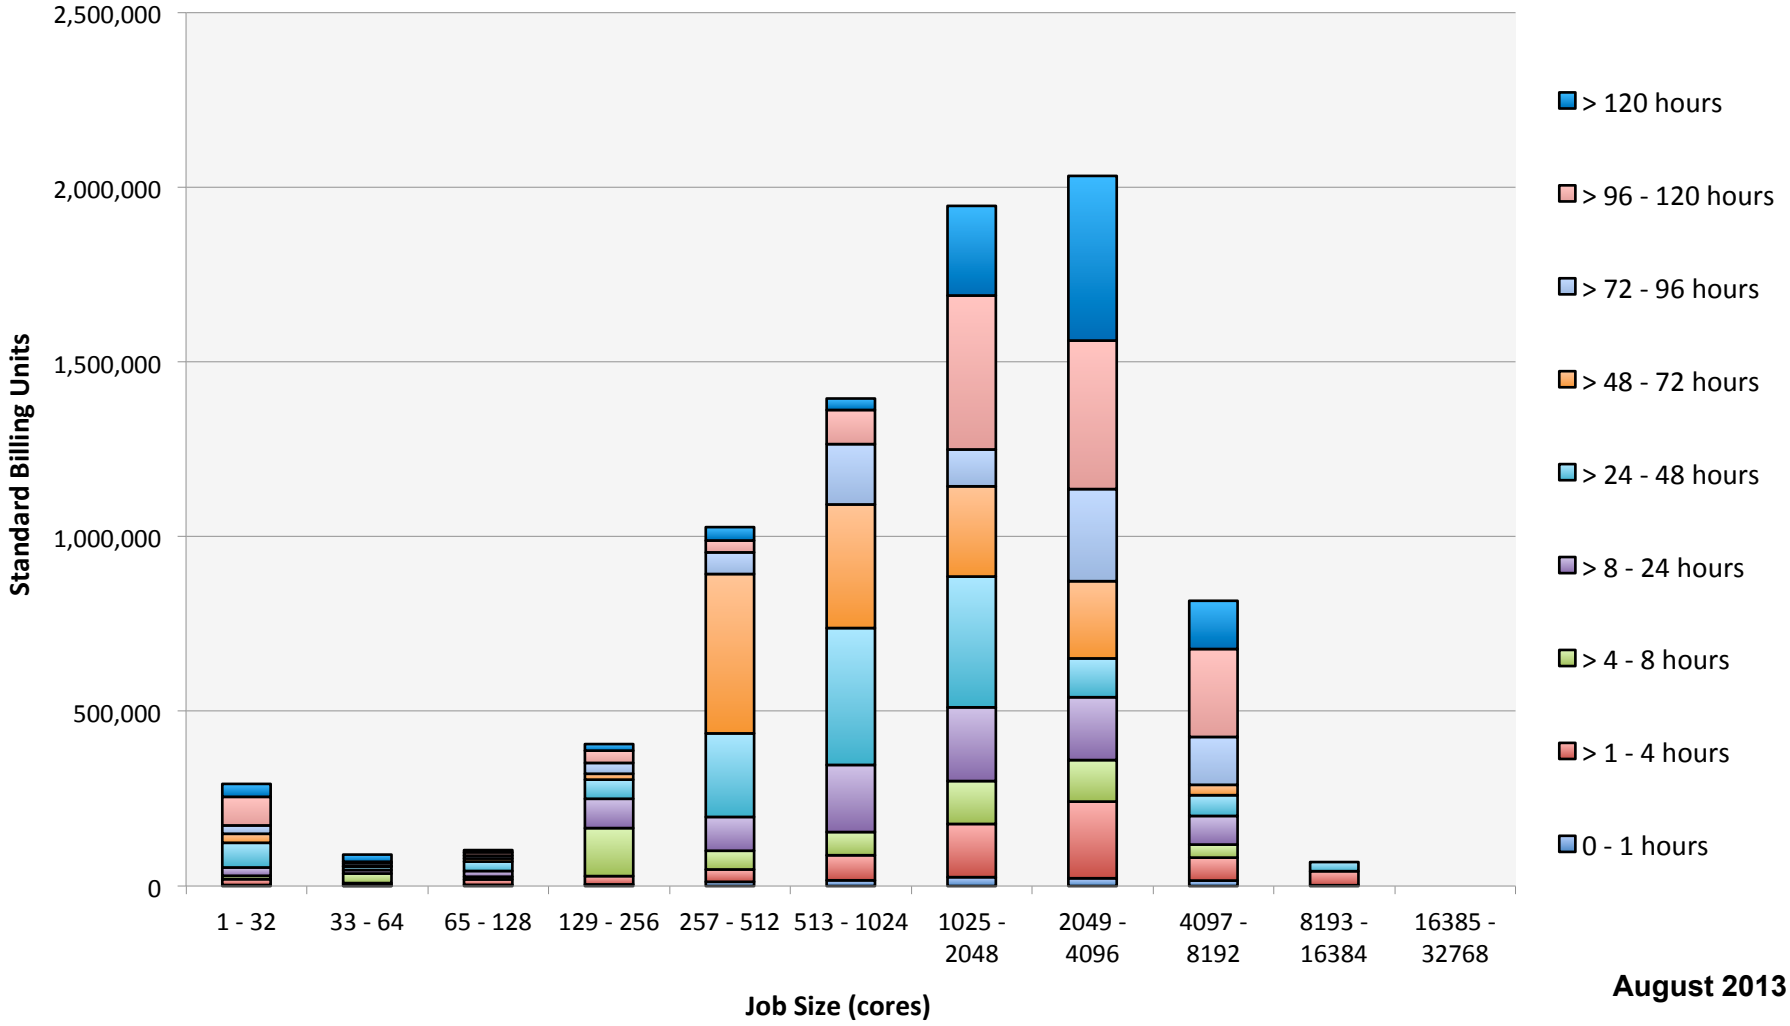

Endeavour:
SBUs Reported, Normalized to 30-Day Month

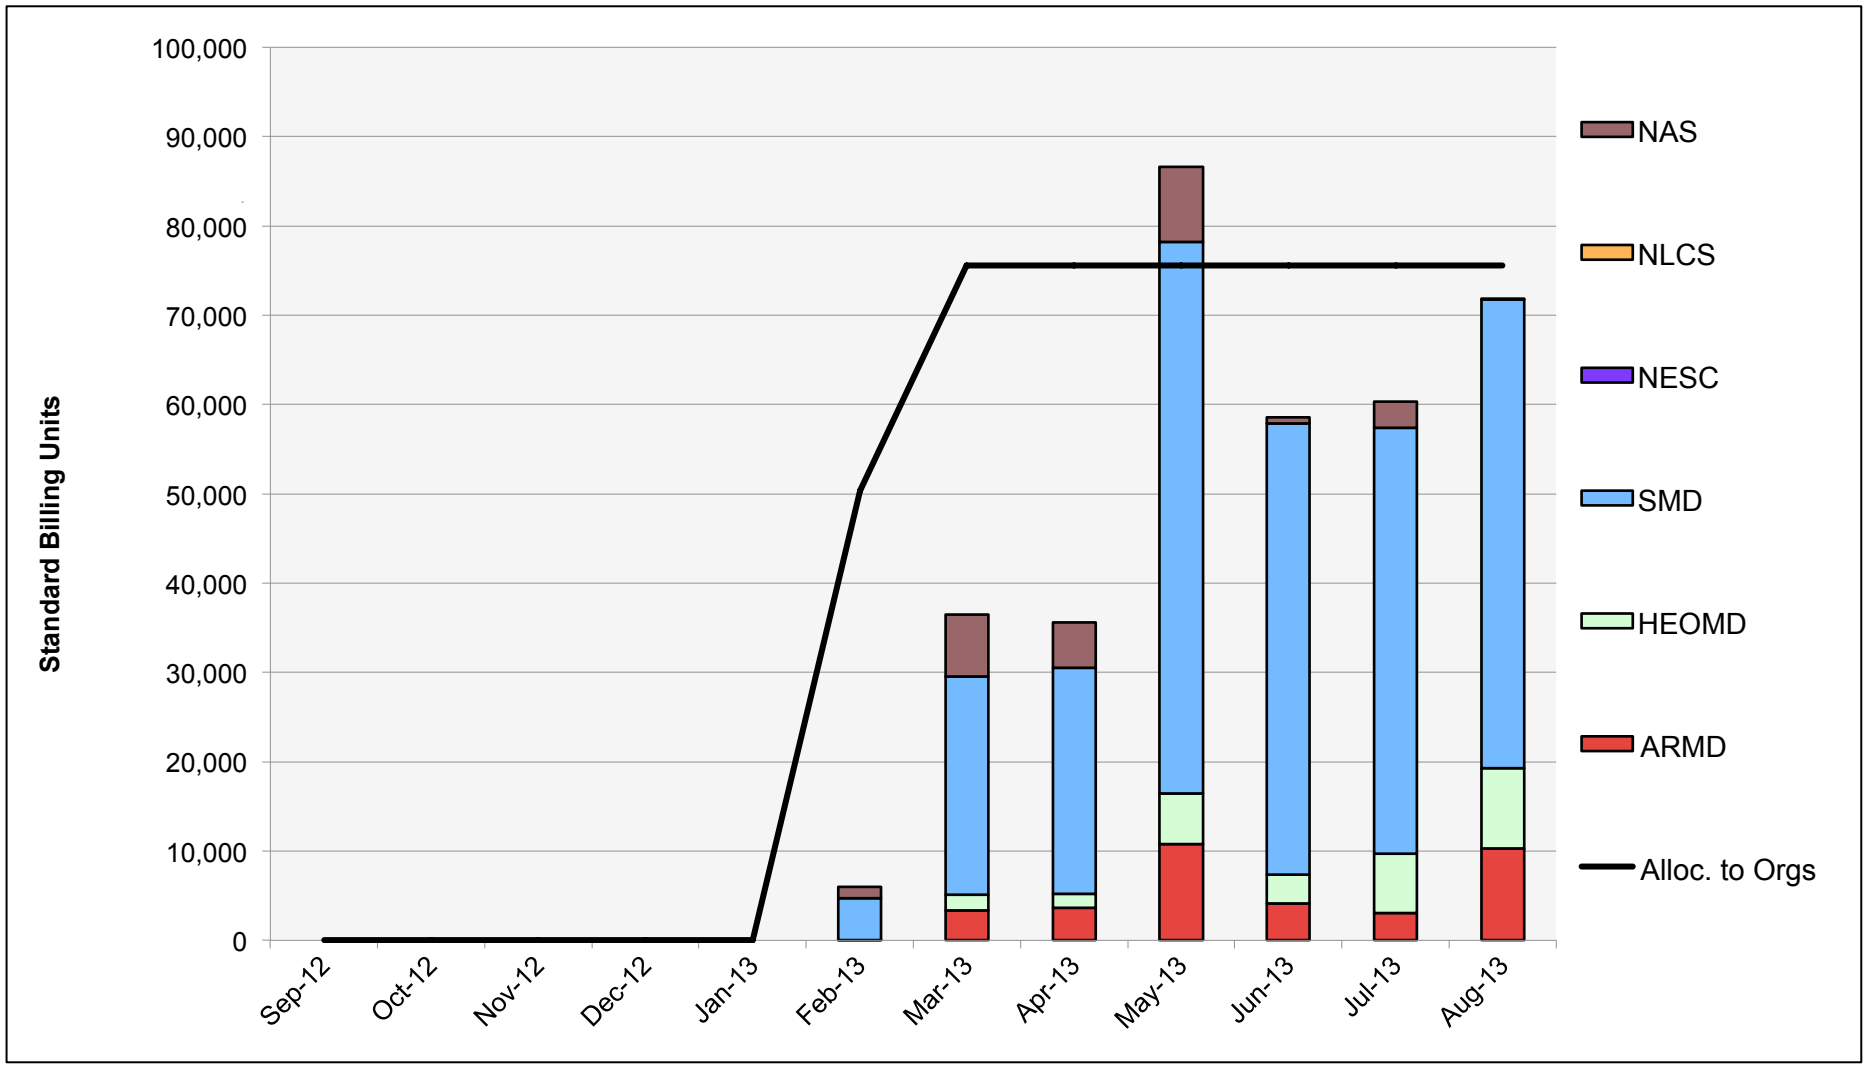

Endeavour:
Monthly Utilization by Job Length

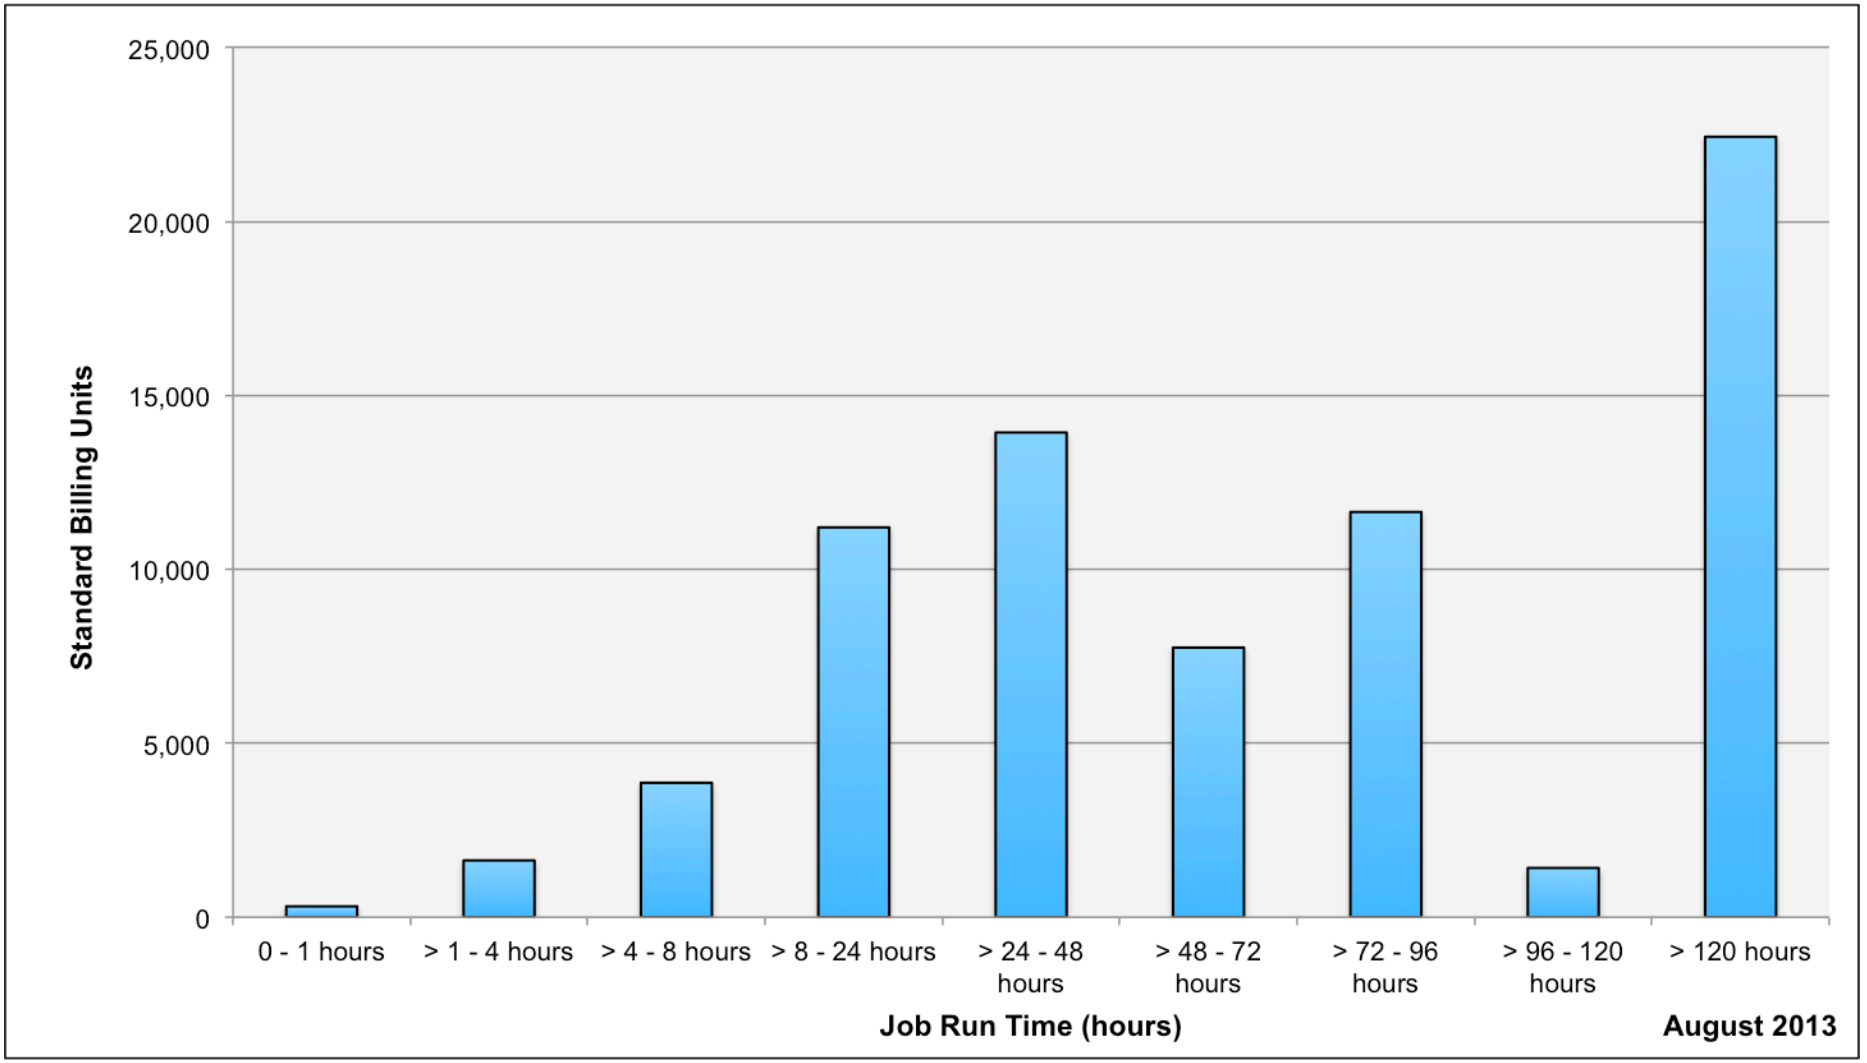

Endeavour:
Monthly Utilization by Size and Length

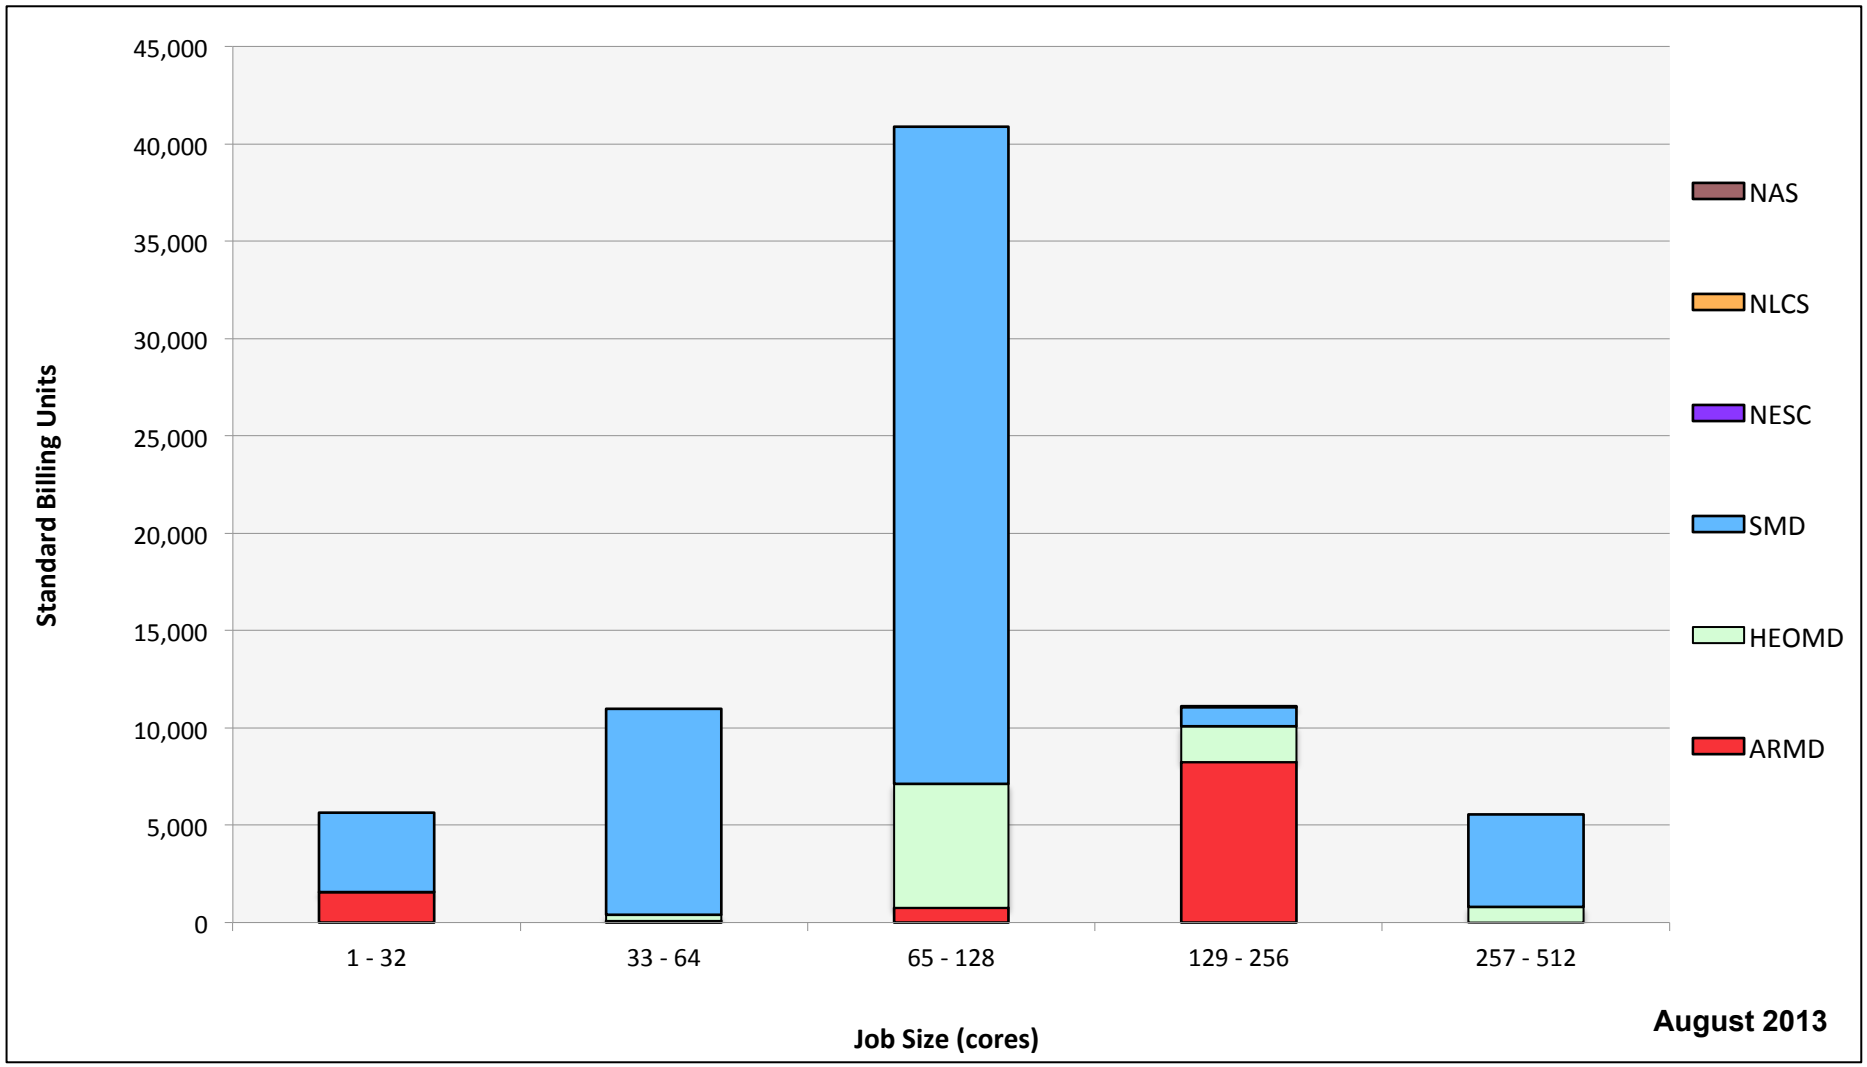

Endeavour:
Monthly Utilization by Size and Mission